INTELLIGENCE BRIEFING
Security Command Center
TLP:CLEAR
2026-06-25 06:53 UTC

EO 14409 Puts AI Security on a 60-Day Clock, But Accountability Gaps Remain

GOVERNANCE | MEDIUM | CVSS 5.0

SCC Item ID	SCC-GOV-2026-0078
Type	Governance
Severity	MEDIUM
CVSS Base Score	5.0
Affected Products	Federal civilian .gov infrastructure, CISA federal systems, DHS/Commerce/NSA agency environments; CrowdStrike Falcon Platform and Charlotte AI referenced as relevant vendor tooling
Discovery Source	Rss:T1 Threatintel

Executive Summary

President Trump signed Executive Order 14409 on June 2, 2026, mandating federal agencies harden systems against AI-enabled threats within 30-60 days, with coordination across DHS, Commerce, NSA, and CISA. Federal civilian infrastructure and agencies operating AI-integrated systems face immediate compliance deadlines, while the EO's classified benchmarking framework and voluntary participation model for non-federal entities create significant enforceability gaps. For organizations with federal contracts or critical infrastructure designations, the absence of enforceable non-federal requirements does not eliminate risk exposure; it shifts accountability to internal governance while AI-enabled adversaries continue accelerating attack volume.

Technical Analysis

Executive Order 14409 establishes a 30-60 day timeline for federal agencies to implement AI security hardening across civilian .gov infrastructure. Key technical mandates include a federal vulnerability clearinghouse for AI-related weaknesses, an interagency coordination structure spanning DHS, Commerce, NSA, and CISA, and a benchmarking framework (details classified). Three CWE-mapped risk surfaces are central to the EO's operational scope: CWE-287 (Improper Authentication) reflecting machine identity gaps in agentic AI deployments, CWE-1395 (Dependency on Vulnerable Third-Party Component) reflecting open-source supply chain exposure in AI toolchains, and CWE-693 (Protection Mechanism Failure) reflecting shadow AI deployments operating without security guardrails. MITRE ATT&CK techniques in scope include T1588.006 (AI-accelerated vulnerability discovery), T1195 (supply chain compromise targeting open-source repositories), T1078 (valid account abuse via machine identity in agentic systems), T1550 (use of alternate authentication

material), and T1584 (infrastructure compromise). CrowdStrike's 2025 threat report documents an 89% year-over-year increase in AI-enabled attack volume, providing operational context for the EO's hardening mandate. No CVE is associated; this is a governance and policy development item with direct operational implications.

Action Checklist

1. Step 1: Scoping. Federal agencies: map all AI-integrated systems against the 30-60 day EO 14409 timeline and identify gaps in machine identity management (CWE-287) across agentic deployments. Reference NIST AC-2 (Account Management) to enumerate AI service accounts and non-human identities. Federal contractors and critical infrastructure operators should conduct equivalent scoping now, ahead of any future mandatory extension.
2. Step 2: Detection. Audit for shadow AI deployments (CWE-693) by reviewing software inventories (CIS 2.1) and network traffic for unauthorized AI service calls. Query authentication logs for machine-to-machine sessions lacking MFA enforcement (CIS 6.5). Review open-source dependencies in AI toolchains for vulnerable third-party components (CWE-1395) using CIS 2.2 (Ensure Authorized Software is Currently Supported).
3. Step 3: Eradication. Enforce least-privilege access controls on all AI service accounts per NIST AC-6 (Least Privilege). Remove or formally document all shadow AI tools identified in Step 2; apply CIS 2.3 (Address Unauthorized Software) procedures. Rotate credentials for machine identities with excessive privilege per NIST IA-5 (Authentication and Authorization).
4. Step 4: Recovery. Validate that all AI-integrated services are included in audit logging per NIST AU-2 (Event Logging) and AU-12 (Audit Record Generation). Confirm MFA enforcement on all administrative and externally-exposed AI interfaces per CIS 6.3 and CIS 6.5. Monitor for T1078 and T1550 activity patterns post-remediation using SIEM alerting on machine identity authentication anomalies.
5. Step 5: Post-Incident. Document identified gaps against NIST AC-2, AC-6, and AU-2 for inclusion in your next risk assessment cycle. Establish a process for continuous inventory of AI components (CIS 2.1) and their supply chain dependencies (CWE-1395). Engage with the federal vulnerability clearinghouse once operational to align internal benchmarking with EO 14409 requirements.

IR / Forensic Enrichment

Triage Priority	URGENT
Escalation Criteria	Escalate to CISO and legal counsel immediately if shadow AI deployments identified in Step 2 are found to have processed federal data outside approved system boundaries, or if any AI service account shows authentication activity inconsistent with known integration patterns — either condition may trigger federal incident reporting obligations under FISMA and could constitute a material compliance failure under EO 14409's 30-day deadline.

Recovery Notes	Post-remediation monitoring should remain elevated for a minimum of 30 days to align with the EO 14409 compliance window, focusing specifically on machine identity authentication anomalies from AI service accounts modified during eradication. Verify that CrowdStrike Falcon Platform telemetry and Charlotte AI connector logs are feeding into your audit log pipeline and that no AI-integrated service was inadvertently excluded from AU-2 logging scope during the credential rotation in Step 3. Re-run the software inventory diff and IAM permission audit at Day 15 and Day 30 to confirm no shadow AI deployments have re-emerged and that scope-restricted service accounts have not been re-elevated.
Forensic Artifacts	Cloud provider IAM audit logs (AWS CloudTrail `LookupEvents` filtered on AI service principal ARNs, Azure Activity Log filtered on service principal object IDs, GCP Cloud Audit Logs for AI Platform service accounts) — these capture the exact scope of machine identity activity during the EO 14409 compliance window and are the primary evidence source for overprivileged agentic deployments | CrowdStrike Falcon Platform API audit logs and Charlotte AI connector authentication records — specific to the vendor tooling named in the advisory, these reveal whether AI-integrated detection capabilities were accessed by unauthorized machine identities or operated outside approved configuration boundaries | Windows Security Event Log Event ID 4648 (Explicit Credential Logon) and 4624 Type 5 (Service Logon) filtered on AI service account SIDs — documents machine-to-machine authentication patterns that lack MFA enforcement, directly evidencing the gap identified in Step 2 | Software inventory snapshots from osquery (`SELECT name, version, source, install_time FROM programs`) and pip/npm dependency manifests from AI toolchain hosts — establishes the pre-remediation state of shadow AI deployments and open-source supply chain components for CIS 2.1/2.2 gap documentation | DNS query logs and network flow records showing outbound connections to AI provider API endpoints (api.openai.com, api.anthropic.com, bedrock.us-east-1.amazonaws.com, and equivalent Azure/GCP AI service FQDNs) — identifies unauthorized AI service calls not reflected in the approved software inventory, supporting the shadow AI detection analysis in Step 2

Per-Action IR Details

Step 1: Scoping — Federal agencies: map all AI-integrated systems against the 30–60 day EO 14409 timeline and identify gaps in machine identity management (CWE-287) across agentic deployments. Reference NIST AC-2 (Account Management) to enumerate AI service accounts and non-human identities. Federal contractors and critical infrastructure operators should conduct equivalent scoping now, ahead of any future mandatory extension.

NIST Phase: Preparation

Reference: NIST 800-61r3 §2 — Preparation: Establishing IR capability, asset inventories, and pre-incident tooling before adverse events occur

Controls: NIST AC-2 (Account Management), CIS 1.1 (Establish and Maintain Detailed Enterprise Asset Inventory), CIS 2.1 (Establish and Maintain a Software Inventory), CIS 5.1 (Establish and Maintain an Inventory of Accounts)

Compensating: Run `net user /domain` and `Get-ADServiceAccount -Filter *` (Windows/AD) or `aws iam list-users` / `gcloud iam service-accounts list` to enumerate machine identities. Cross-reference against a manually maintained spreadsheet of AI-integrated services (e.g., CrowdStrike Charlotte AI API keys, OpenAI org tokens, Azure AI service principals). A 2-person team can divide cloud vs. on-prem enumeration and merge results in a shared CSV. Use osquery with `SELECT * FROM users WHERE type='special';` on Linux endpoints to surface service accounts missed by AD queries.

Evidence: This is a preparation/scoping step and does not alter live system state, so order-of-volatility sequencing is not triggered. However, document the current state of AI service account configurations before any changes: export IAM role assignments, service principal permission scopes, and API key metadata (creation date, last-used timestamp) from all cloud providers. Capture CrowdStrike Falcon Platform API client configurations and Charlotte AI integration

connector settings as baseline evidence for gap analysis. These records establish the pre-EO 14409 state for audit purposes.

Step 2: Detection — Audit for shadow AI deployments (CWE-693) by reviewing software inventories (CIS 2.1) and network traffic for unauthorized AI service calls. Query authentication logs for machine-to-machine sessions lacking MFA enforcement (CIS 6.5). Review open-source dependencies in AI toolchains for vulnerable third-party components (CWE-1395) using CIS 2.2 (Ensure Authorized Software is Currently Supported).

NIST Phase: Detection Analysis

Reference: NIST 800-61r3 §3.2 — Detection and Analysis: Monitoring for indicators of unauthorized activity, correlating authentication logs, and analyzing software inventory anomalies

Controls: NIST AU-2 (Event Logging), NIST AU-6 (Audit Record Review, Analysis, And Reporting), CIS 2.1 (Establish and Maintain a Software Inventory), CIS 2.2 (Ensure Authorized Software is Currently Supported), CIS 6.5 (Require MFA for Administrative Access), CIS 8.2 (Collect Audit Logs)

Compensating: Use Wireshark or `tcpdump -i any -w ai_traffic.pcap 'host api.openai.com or host api.anthropic.com or host falcon.crowdstrike.com'` to capture outbound AI API calls not in your approved inventory. Parse DNS query logs with `grep -E 'openai|anthropic|bedrock|vertex|charlotte' /var/log/named/queries.log` to surface unauthorized AI service resolutions. For authentication log review without a SIEM, run `Get-WinEvent -LogName Security | Where-Object {$_.Id -eq 4648}` (explicit credential logon) and filter for service account names to find machine-to-machine sessions. Run `pip-audit` or `npm audit` against AI toolchain dependency trees to identify vulnerable open-source components.

Evidence: Before acting on any findings, capture: (1) raw authentication log exports from Windows Security Event Log (Event IDs 4624, 4648, 4776) and Linux `/var/log/auth.log` or `/var/log/secure` filtered on AI service account names; (2) a full network flow snapshot showing connections to known AI provider IP ranges (OpenAI ASN 54113, AWS Bedrock endpoints, Azure OpenAI endpoints) to establish a baseline of shadow AI call volume; (3) the current software inventory state from each endpoint via osquery `SELECT name, version, install_time FROM programs;` before any removal actions in Step 3 overwrite this evidence.

Step 3: Eradication — Enforce least-privilege access controls on all AI service accounts per NIST AC-6 (Least Privilege) and apply D3-UAP (User Account Permissions) to restrict agentic identity scope. Remove or formally document all shadow AI tools identified in Step 2; apply CIS 2.3 (Address Unauthorized Software) procedures. Rotate credentials for machine identities with excessive privilege per D3-CRO (Credential Rotation).

NIST Phase: Eradication

Reference: NIST 800-61r3 §3.4 — Eradication: Removing unauthorized components from the environment and hardening identity configurations to eliminate the threat foothold

Controls: NIST AC-6 (Least Privilege), NIST AC-2 (Account Management), CIS 2.3 (Address Unauthorized Software), CIS 5.4 (Restrict Administrator Privileges to Dedicated Administrator Accounts)

Compensating: Use `Remove-AppxPackage` (Windows) or `apt purge` / `yum remove` for unauthorized AI client software identified in Step 2. Scope-restrict AI service principals via `az ad sp update --id --set 'appRoles=[]'` (Azure) or IAM inline policy replacement on AWS. For credential rotation without an enterprise PAM tool, use a scripted sequence: (1) generate new API key via provider console, (2) update the secret in the consuming application's config, (3) revoke the old key, (4) validate with a test call — document each step with timestamps for EO 14409 audit evidence. Do NOT rotate credentials before capturing the artifacts listed in the evidence field below.

Evidence: ORDER OF VOLATILITY — capture BEFORE revoking sessions or rotating credentials: (1) Export all active API key last-used timestamps and associated source IPs from CrowdStrike Falcon API audit logs and Charlotte AI connector logs — this establishes whether overprivileged machine identities were actively exploited before remediation; (2) Capture `netstat -ano` or `Get-NetTCPConnection` output on hosts running agentic AI workloads to record active sessions that will be severed by credential rotation; (3) Export the current IAM/RBAC permission set for each AI service account as a JSON snapshot before scope restriction — this is the forensic baseline proving the pre-remediation permission state for EO 14409 compliance documentation.

Step 4: Recovery — Validate that all AI-integrated services are included in audit logging per NIST AU-2 (Event Logging) and AU-12 (Audit Record Generation). Confirm MFA enforcement on all administrative and externally-exposed AI interfaces per CIS 6.3 and CIS 6.5. Monitor for T1078 and T1550 activity patterns post-remediation using SIEM alerting on machine identity authentication anomalies.

NIST Phase: Recovery

Reference: NIST 800-61r3 §3.5 — Recovery: Restoring systems to verified operational state, confirming integrity of security controls, and establishing post-remediation monitoring

Controls: NIST AU-2 (Event Logging), NIST AU-12 (Audit Record Generation), NIST AU-6 (Audit Record Review, Analysis, And Reporting), CIS 6.3 (Require MFA for Externally-Exposed Applications), CIS 6.5 (Require MFA for Administrative Access), CIS 8.2 (Collect Audit Logs)

Compensating: Without a SIEM, deploy Sigma rules converted to native query syntax: use `Get-WinEvent -FilterHashtable @{LogName='Security'; Id=4625,4648,4776}` scheduled as a daily PowerShell task and output to a reviewed CSV to detect credential stuffing against AI service accounts (Valid Accounts and Use Alternate Authentication Material patterns). Validate logging coverage by running `auditpol /get /category:*` (Windows) and confirming AI service process activity is captured. For MFA validation on externally exposed AI interfaces, manually test each admin login flow and document the MFA challenge response as evidence for EO 14409 auditors. Note: ATT&CK technique IDs T1078 and T1550 referenced in the source step identify the attacker behavior to detect — the defensive controls above govern the monitoring posture.

Evidence: This step validates controls rather than altering active threat state, so primary order-of-volatility risk is lower. However, before enabling new logging pipelines that may overwrite existing log buffers, export current log state: (1) Windows Security Event Log export filtered on the AI service account SIDs identified in Steps 1–3, covering the full EO 14409 scoping window; (2) CrowdStrike Falcon Platform detection telemetry for any Charlotte AI connector authentication events during the remediation period; (3) Cloud provider audit trail exports (AWS CloudTrail, Azure Activity Log, GCP Cloud Audit Logs) for all AI service principal activity since the Step 2 detection window — these form the evidentiary chain for post-incident reporting.

Step 5: Post-Incident — Document identified gaps against NIST AC-2, AC-6, and AU-2 for inclusion in your next risk assessment cycle. Establish a process for continuous inventory of AI components (CIS 2.1) and their supply chain dependencies (CWE-1395). Engage with the federal vulnerability clearinghouse once operational to align internal benchmarking with EO 14409 requirements.

NIST Phase: Post Incident

Reference: NIST 800-61r3 §4 — Post-Incident Activity: Lessons learned documentation, process improvement, policy updates, and intelligence sharing to improve future detection and response posture

Controls: NIST AC-2 (Account Management), NIST AC-6 (Least Privilege), NIST AU-2 (Event Logging), CIS 2.1 (Establish and Maintain a Software Inventory), CIS 2.2 (Ensure Authorized Software is Currently Supported), CIS 7.1 (Establish and Maintain a Vulnerability Management Process), CIS 7.2 (Establish and Maintain a Remediation Process)

Compensating: Document gaps in a structured after-action report template mapped to the EO 14409 control domains (machine identity management, agentic deployment scope, AI supply chain visibility). Use a free GRC spreadsheet (CISA's CPGs self-assessment template is publicly available) to track AC-2, AC-6, and AU-2 gap closure. For continuous AI component inventory, implement a monthly cron job running `pip list --format=json > ai_deps_$(date +%F).json` on all AI workload hosts and diff against the prior month's output to detect supply chain drift. Assign one team member as the EO 14409 clearinghouse liaison once the federal vulnerability clearinghouse is operationalized.

Evidence: No live system state is altered in this phase. Consolidate the following as the final forensic record for the EO 14409 compliance file: (1) the pre-remediation AI service account permission snapshots captured in Step 3; (2) the authentication log exports from Step 4 covering the full incident window; (3) the shadow AI software inventory from Step 2 showing what was found, when, and its disposition (removed vs. formally documented exception); (4) timestamped evidence of MFA enforcement validation from Step 4. This package constitutes the audit trail for CISA federal system oversight and supports any future DHS/NSA benchmarking review under EO 14409's classified framework.

Detection Guidance

Focus detection efforts on three surfaces mapped to the EO's risk framework. For machine identity gaps (CWE-287, T1078): query authentication logs for service accounts or API keys authenticating without MFA, especially those with broad permissions in AI orchestration layers; flag sessions using T1550 patterns (pass-the-token, alternate auth material). For shadow AI (CWE-693): compare active network connections and DNS queries against your authorized software inventory (CIS 2.1); outbound calls to AI API endpoints (e.g., OpenAI, Anthropic, Hugging Face) not in the approved software list indicate unmanaged deployments. Monitor for newly created service accounts associated with AI tooling. For supply chain exposure (CWE-1395, T1195): audit open-source package manifests in AI/ML pipelines for dependencies flagged in CISA advisories or the forthcoming federal vulnerability clearinghouse. Reference NIST AU-6 (Audit Record Review, Analysis, and Reporting) for ongoing log review cadence. No IOCs are associated with this governance item; detection is posture-based rather than indicator-based.

Framework Mappings

MITRE-ATTACK

- **T1588.006** — Vulnerabilities
- **T1195** — Supply Chain Compromise
- **T1584** — Compromise Infrastructure
- **T1078** — Valid Accounts
- **T1588** — Obtain Capabilities
- **T1550** — Use Alternate Authentication Material

NIST-800-53R5

- **SA-9** — External System Services
- **SR-2** — Supply Chain Risk Management Plan
- **SR-3** — Supply Chain Controls and Processes
- **SI-7** — Software, Firmware, and Information Integrity
- **AC-2** — Account Management
- **AC-6** — Least Privilege
- **IA-2** — Identification and Authentication (Organizational Users)
- **IA-5** — Authenticator Management
- **IA-8** — Identification and Authentication (Non-Organizational Users)
- **IR-5** — Incident Monitoring

OWASP-TOP10-2021

- **A07:2021** — Identification and Authentication Failures

CIS-V8

- **6.3** — Require MFA for Externally-Exposed Applications
- **6.4** — Require MFA for Remote Network Access

- **6.5** — Require MFA for Administrative Access
- **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(d)** — Person or Entity Authentication

ISO-27001-2022

- **A.8.8** — Management of technical vulnerabilities
- **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
T1588.006	Vulnerabilities	Resource-Development
T1195	Supply Chain Compromise	Initial-Access
T1584	Compromise Infrastructure	Resource-Development
T1078	Valid Accounts	Defense-Evasion
T1588	Obtain Capabilities	Resource-Development
T1550	Use Alternate Authentication Material	Defense-Evasion

Sources

Source	URL	Tier
Blog	https://www.crowdstrike.com/en-us/blog/after-executive-order-14409-...	T3
	https://industrialcyber.co/regulation-standards-and-compliance/trum...	T3
	https://techpolicy.press/transparency-and-accountability-gaps-in-tr...	T3
	https://www.hstoday.us/subject-matter-areas/cybersecurity/white-hou...	T3

Source	URL	Tier
Cybersecurity For The Federal Government - CrowdStrike	https://www.crowdstrike.com/en-us/solutions/federal-government/	T3

Generated 2026-06-25 06:53 UTC by TJS Security Command Center