SPECTRA — APRIL 2026

Security Policy, Emerging Cyber Threats, Research & AI

A monthly digest of cybersecurity policy, standards, threats, and AI developments relevant to DoD and federal practitioners. Compiled from public government and industry sources.

Executive Summary

Iranian-Affiliated APT Actors Exploit PLCs Across US Critical Infrastructure — Iran-affiliated APT actors are actively exploiting internet-facing OT devices including Rockwell Automation/Allen-Bradley PLCs, causing confirmed disruptions across U.

FBI Disrupts Russian GRU Forest Blizzard Router Espionage Network of 18,000 Devices (Operation Masquerade) — The FBI and international partners disrupted Forest Blizzard (APT28/GRU), which hijacked over 18,000 SOHO routers worldwide by modifying DNS settings to intercept Microsoft account credentials and tokens with no malware deployed, making detection extremely difficult.

Suspected Chinese Breach of FBI Surveillance System Exposed Wiretap Target Phone Numbers — A suspected Chinese state-sponsored intrusion into an FBI network system exposed the phone numbers of individuals under active U.

Pro-Iran Hackers Breach FBI Director Email and Stryker Defense Supplier as Iran War Cyber Campaign Expands — Pro-Iran hackers claimed and apparently executed a breach of the FBI Director email account with leaks assessed as authentic, while a separate Iran-linked group compromised Stryker Corporation -- a major DoD medical device supplier -- causing a global Windows network outage and prompting CISA and FBI formal investigation.

CISA KEV Additions: 15 Critical Vulnerabilities Including Ivanti EPMM, Fortinet FortiClient, Citrix NetScaler, and Cisco FMC — CISA added 15 entries to the Known Exploited Vulnerabilities catalog in late March through April 2026 covering products widespread in federal networks: Ivanti EPMM unauthenticated RCE (CVE-2026-1340, remediation due 4/11); Fortinet FortiClient EMS with emergency patch issued (CVE-2026-35616); Citrix NetScaler out-of-bounds read (CVE-2026-3055); Cisco Secure FMC Java deserialization as root (CVE-2026-20131); F5 BIG-IP RCE (CVE-2025-53521); Aquasecurity Trivy embedded malicious code affecting DevSecOps pipelines (CVE-2026-33634); Langflow AI platform code injection (CVE-2026-33017); and multiple Google Chrome/Chromium browser vulnerabilities.

Policy & Compliance Updates

Trump Administration National Cyber Strategy Emphasizes Offensive Response and Rejects Incremental Approaches

nextgov — https://www.nextgov.com/cybersecurity/2026/03/trumps-new-cyber-strategy-details-more-offensive-response-cyber-threats/411963/

The Trump administration released its National Cyber Strategy explicitly prioritizing offensive cyber operations as a primary response mechanism, directly stating the administration will not apply partial or ambiguous

strategies and will take a more aggressive posture than prior administrations. National Cyber Director Sean Cairncross clarified the government does not envision industry conducting offensive operations independently but expects private sector technical expertise to inform U.S. government offensive and defensive decisions. DoD and IC components should review how the strategy affects cyber operational authorities, interagency coordination structures, and the scope of permissible industry-government cyber collaboration.

White House Expands Offensive Cyber Market as Policy Boundaries for Private Sector Remain Undefined

nextgov — https://www.nextgov.com/cybersecurity/2026/04/us-push-counter-hackers-draws-industry-deeper-offensive-cyber-debate/412770/

The White House is expanding the market for offensive cyber capabilities and increasing private sector involvement in U.S. cyber operations even as the policy boundaries around authorized use of those capabilities remain ambiguous, creating legal and operational risk for industry partners. The administration is actively building out ecosystem support for offensive tools while deferring resolution of the policy questions that would clarify what industry can and cannot do. Federal cyber practitioners should monitor forthcoming policy clarifications affecting the scope of permissible private sector engagement in offensive cyber activity before entering new partnerships.

Trump FY27 Budget Proposes $360 Million Cut to CISA, Eliminates Election Security Program and 860 Positions

nextgov — https://www.nextgov.com/cybersecurity/2026/04/trump-proposes-cutting-cisa-election-security-program-fy27-budget/412672/

The FY27 budget proposal includes approximately $700 million in CISA program cuts resulting in a net $360 million funding reduction, elimination of the election security program, and approximately 860 position cuts -- significantly constraining CISA capacity to support state and local governments at a time of active multi-nation threat operations. The reductions come as CISA is simultaneously engaged in Iranian, Russian, and North Korean threat response while managing several major vulnerability and supply chain incidents. Federal programs relying on CISA shared services, threat information sharing, or election security support should begin contingency planning for reduced CISA availability.

FedRAMP Rev5 Modernization Progress: RFC Outcomes Published on Machine-Readable Packages, Certifications, and Incident Communications

fedramp_notices — https://fedramp.gov/notices/0009

FedRAMP published initial outcomes from three recently closed public comment periods: RFC-0024 (Rev5 Machine-Readable Security Packages), RFC-0023 (Rev5 Program Certifications for 3PAOs and agencies), and RFC-0022 (Leveraging External Frameworks); the program also updated incident communication procedures via RFC-0031 added to the changelog. These outcomes signal continued Rev5 transition progress that will reshape how CSPs document authorizations, how assessors are certified, and how agencies accept external security frameworks as equivalent evidence. Federal cloud acquisition officers and authorizing officials should review published outcomes to plan compliance with forthcoming Rev5 requirements.

FCC Seeks Comment on Banning Previously Authorized Covered Communications Equipment From Chinese Manufacturers

federal_register — https://www.federalregister.gov/documents/2026/04/06/2026-06653/seeking-comment-on-prohibiting-importation-and-marketing-of-previously-authorized-covered

The FCC Public Safety and Homeland Security Bureau is seeking comment on prohibiting continued importation and marketing of certain previously authorized communications equipment from entities on the Covered List -- primarily Chinese manufacturers -- that are now determined to pose unacceptable national security risk, potentially requiring replacement of legacy gear already in service on U.S. networks. The action would extend prior new-equipment bans to gear already deployed, with potential compliance obligations for federal agencies running existing covered equipment. Federal technology officers should assess Covered List equipment in inventories and monitor this rulemaking for potential compliance deadlines.

FERC Order 918: CIP-003-11 Closes Cybersecurity Management Gap for Low-Impact Electric Grid Facilities

federal_register — https://www.federalregister.gov/documents/2026/03/24/2026-05711/order-no-918-critical-infrastructure-protection-reliability-standard-cip-003-11-cyber

FERC Order 918 approved NERC CIP-003-11, adding mandatory cybersecurity management controls for low-impact bulk electric system facilities -- closing a regulatory gap where coordinated attacks on individually low-impact sites could produce aggregate grid-wide disruption. The standard imposes new requirements on utilities previously subject to minimal mandatory controls, including facilities serving military installations. DoD energy officers and facility security managers contracting with utilities should understand how CIP-003-11 affects cyber posture requirements for power providers serving critical federal sites.

FERC Order 919: Virtualization Reliability Standards Close Cybersecurity Gap for Virtualized Electric Grid Control Systems

federal_register — https://www.federalregister.gov/documents/2026/03/24/2026-05716/order-no-919-virtualization-reliability-standards

FERC Order 919 approved 11 modified CIP Reliability Standards and updated 22 NERC Glossary definitions to address cybersecurity risks from virtualization of bulk electric system control environments, closing a significant gap as utilities migrate grid control to virtualized and cloud-based platforms. The standards establish security baselines for virtualized BES cyber systems that previously had no explicit CIP coverage. Federal energy and facility security officers should note these standards affect cybersecurity compliance expectations for utility partners supplying power to DoD and federal installations.

Treasury Launches Cyber Threat Intelligence Sharing Initiative with Cryptocurrency Sector

nextgov — https://www.nextgov.com/cybersecurity/2026/04/treasury-debuts-effort-share-cyber-threat-intel-crypto-firms/412756/

Treasury launched a bilateral cyber threat intelligence sharing program with cryptocurrency firms, treating the crypto sector as part of core national financial infrastructure and acknowledging that North Korean and other state-sponsored actors are aggressively targeting crypto platforms as high-value financial attack surfaces. The initiative expands Treasury sector-specific threat sharing beyond traditional banking into digital asset markets with implications for how fintech risks are managed across financial sector oversight. Federal financial oversight professionals and program managers with fintech acquisition responsibilities should understand Treasury evolving posture on crypto infrastructure security.

Publications & Standards

NIST Cyber AI Profile: Second Workshop Outcomes Inform Next Draft Integrating CSF and AI RMF

nist_csrc — https://www.nist.gov/blogs/cybersecurity-insights/reflections-second-nist-cyber-ai-profile-workshop

NIST published reflections from its second Cyber AI Profile workshop, summarizing community input on the Preliminary Draft guidance integrating the Cybersecurity Framework and AI Risk Management Framework to support organizations using AI for cybersecurity, defending against AI-enabled attacks, and protecting AI systems. The profile is expected to become a reference document for federal AI security requirements and will likely inform future OMB and CISA guidance on securing AI in government environments. Federal agencies deploying AI systems for cybersecurity use cases should engage with the ongoing development process and treat this profile as a likely compliance baseline.

NIST SP 800-63 Revision 4 Final: Updated Digital Identity Guidelines Now Published

nist_csrc — https://www.nist.gov/blogs/cybersecurity-insights/lets-get-digital-updated-digital-identity-guidelines-are-here

NIST published the final Revision 4 of SP 800-63 Digital Identity Guidelines, completing a four-year process incorporating approximately 6,000 public comments and updating guidance on identity proofing, authentication, and federation including updated phishing-resistant MFA requirements and identity assurance level definitions relevant to Zero Trust implementation. Rev 4 is the authoritative guidance for federal ICAM programs and a foundational ZTA compliance reference. Federal agencies should immediately assess compliance gaps against Rev 4 requirements and develop remediation plans, particularly for identity proofing processes and authenticator lifecycle management.

NIST IoT Cybersecurity Future Directions Workshop Shapes Revision of NIST IR 8259 Manufacturer Guidance

nist_csrc — https://www.nist.gov/blogs/cybersecurity-insights/all-aboard-nist-cybersecurity-iot-program-headed-our-next-stop-share

NIST Cybersecurity for IoT Program held its Future Directions Workshop to gather stakeholder input for revising NIST IR 8259 (Foundational Cybersecurity Activities for IoT Device Manufacturers), the guidance underlying federal IoT procurement requirements under the IoT Cybersecurity Improvement Act. Updates to IR 8259 will directly affect baseline cybersecurity standards applied to IoT acquisitions across the federal government, critical infrastructure, and the defense industrial base. Program managers responsible for IoT acquisitions and supply chain risk management should engage with the public comment process to ensure federal operational needs shape the updated standards.

OWASP GenAI Security Project Updated: 21 AI Risks Catalogued with Separate Frameworks for GenAI and Agentic Systems

dark_reading — https://www.darkreading.com/application-security/owasp-genai-security-project-update-matrix

OWASP released a major update to its GenAI Security Project recognizing 21 distinct generative AI security risks and publishing a tools matrix that separately addresses defense of GenAI systems and agentic AI systems, reflecting their distinct attack surfaces including prompt injection, model inversion, and agent trust boundary violations. This is currently the most comprehensive practitioner-level AI security reference available and is designed to complement NIST AI RMF guidance for federal implementations. Federal agencies deploying GenAI or AI agents should use the OWASP matrix to structure AI security risk assessments and as a basis for procurement security requirements.

Threats & Incidents

Iranian-Affiliated APT Actors Exploit PLCs Across US Critical Infrastructure

cisa_alerts — https://www.cisa.gov/news-events/cybersecurity-advisories/aa26-097a

Iran-affiliated APT actors are actively exploiting internet-facing OT devices including Rockwell Automation/Allen-Bradley PLCs, causing confirmed disruptions across U.S. energy, water, and government facilities following U.S.-Israel strikes against Iran. The joint CISA/FBI/NSA advisory confirms actual PLC disruptions have occurred and that approximately 3,900 devices remain exposed across energy, water, and government sectors. Federal agencies overseeing or operating ICS/SCADA environments should immediately audit internet-facing OT device exposure and apply advisory mitigations as a priority action.

FBI Disrupts Russian GRU Forest Blizzard Router Espionage Network of 18,000 Devices (Operation Masquerade)

cyberscoop — https://cyberscoop.com/fbi-operation-masquerade-russian-gru-router-takedown-brett-leatherman/

The FBI and international partners disrupted Forest Blizzard (APT28/GRU), which hijacked over 18,000 SOHO routers worldwide by modifying DNS settings to intercept Microsoft account credentials and tokens with no malware deployed, making detection extremely difficult. FBI cyber chief Brett Leatherman confirmed unique network propagation risk beyond the routers; CISA separately warned that Russian Intelligence Services are targeting commercial messaging application accounts of U.S. government officials, military personnel, and journalists. Federal agencies should treat SOHO router hygiene as a top credential protection priority and audit all edge devices for unauthorized DNS modification.

Suspected Chinese Breach of FBI Surveillance System Exposed Wiretap Target Phone Numbers

nextgov — https://www.nextgov.com/cybersecurity/2026/04/suspected-chinese-breach-fbi-system-exposed-surveillance-targets-phone-numbers/412612/

A suspected Chinese state-sponsored intrusion into an FBI network system exposed the phone numbers of individuals under active U.S. surveillance, potentially allowing foreign adversaries to identify who the U.S. government is monitoring in ongoing counterintelligence operations. The breach extends a pattern of Chinese targeting of U.S. wiretap and lawful intercept infrastructure representing a severe counterintelligence compromise with direct operational consequences. Agencies conducting or supporting surveillance activities should urgently review access controls, compartmentation, and monitoring protocols for sensitive investigative systems.

Pro-Iran Hackers Breach FBI Director Email and Stryker Defense Supplier as Iran War Cyber Campaign Expands

nextgov — https://www.nextgov.com/cybersecurity/2026/03/pro-iran-hackers-claim-breach-fbi-directors-email/412440/

Pro-Iran hackers claimed and apparently executed a breach of the FBI Director email account with leaks assessed as authentic, while a separate Iran-linked group compromised Stryker Corporation -- a major DoD medical device supplier -- causing a global Windows network outage and prompting CISA and FBI formal investigation. The incidents reflect a deliberate escalation of Iranian offensive cyber operations against senior U.S. officials and defense industrial base companies following U.S.-Israel strikes against Iran. Agencies and DIB contractors should elevate threat monitoring for Iran-affiliated actors and ensure incident reporting channels to CISA and sector ISACs are activated.

CISA KEV Additions: 15 Critical Vulnerabilities Including Ivanti EPMM, Fortinet FortiClient, Citrix NetScaler, and Cisco FMC

cisa_kev — https://www.cisa.gov/known-exploited-vulnerabilities-catalog

CISA added 15 entries to the Known Exploited Vulnerabilities catalog in late March through April 2026 covering products widespread in federal networks: Ivanti EPMM unauthenticated RCE (CVE-2026-1340, remediation due 4/11); Fortinet FortiClient EMS with emergency patch issued (CVE-2026-35616); Citrix NetScaler out-of-bounds read (CVE-2026-3055); Cisco Secure FMC Java deserialization as root (CVE-2026-20131); F5 BIG-IP RCE (CVE-2025-53521); Aquasecurity Trivy embedded malicious code affecting DevSecOps pipelines (CVE-2026-33634); Langflow AI platform code injection (CVE-2026-33017); and multiple Google Chrome/Chromium browser vulnerabilities. Per BOD 22-01, federal agencies must verify patch status and remediate all listed products on required deadlines -- several of which have already passed for agencies that have not acted.

North Korean Actors Compromise Axios NPM Package in Precision Open-Source Supply Chain Attack

nextgov — https://www.nextgov.com/cybersecurity/2026/03/north-korea-linked-hackers-suspected-axios-open-source-hijack-google-analysts-say/412523/

Google analysts attributed the brief compromise of the Axios NPM JavaScript package -- used in millions of projects including government and enterprise applications -- to North Korean threat actors who used precision social engineering against the package maintainer, consistent with established DPRK supply chain tactics. Google warned the full breadth of impact remains unclear and will likely be far-reaching across the software ecosystem. Federal software development teams and acquisition officials should audit Axios dependency versions across all managed software and review software composition analysis and SBOM processes.

BlueHammer Windows Zero-Day Exploit Publicly Released, Exposing Microsoft Vulnerability Disclosure Failures

dark_reading — https://www.darkreading.com/vulnerabilities-threats/bluehammer-windows-exploit-microsoft-bug-disclosure-issues

A researcher publicly released a proof-of-concept Windows zero-day exploit enabling local privilege escalation to full system takeover under the alias Chaotic Eclipse, citing an unresolved dispute with Microsoft over coordinated disclosure practices. The release immediately exposes unpatched Windows systems to exploitation while an official Microsoft patch has not been confirmed. Federal IT administrators should monitor MSRC advisories for emergency patching and apply any forthcoming updates on priority timelines.

Storm-1175 Deploys Medusa Ransomware at High Velocity Using N-Day and Zero-Day Exploits

dark_reading — https://www.darkreading.com/threat-intelligence/storm-1175-medusa-ransomware-high-velocity

Microsoft identified Storm-1175, a financially motivated cybercrime group, deploying Medusa ransomware at high velocity by rapidly exploiting both known and zero-day vulnerabilities with speed as the primary operational tactic. The group compresses remediation windows to days or less, making aggressive patching discipline and continuous vulnerability scanning the primary defensive mechanisms. Federal agencies should treat outstanding vulnerability backlogs as high-priority risk given Storm-1175 demonstrated ability to weaponize newly disclosed flaws within days of release.

FBI IC3 Report: Cybercrime Losses Rose 26% to $20.9 Billion in 2025; AI-Fueled Government Impersonation Doubled

cyberscoop — https://cyberscoop.com/fbi-internet-crime-complaint-center-annual-cybercrime-report/

The FBI annual IC3 Internet Crime Report documented $20.9 billion in cybercrime losses in 2025 -- a 26% increase over 2024 -- with government impersonation scam complaints doubling and AI-fueled fraud among the primary growth drivers. The underreporting problem means actual losses are substantially higher than recorded figures. Federal security awareness programs should incorporate AI-enabled government impersonation scenarios into user training and establish updated identity verification protocols for government communications.

AI & Agentic Developments

Anthropic Project Glasswing AI Exploit Tool Raises Federal Cyber Operations Questions; Pentagon Supply Chain Designation Blocked by Court

nextgov — https://www.nextgov.com/cybersecurity/2026/04/anthropics-glasswing-initiative-raises-questions-us-cyber-operations/412721/

Anthropic Project Glasswing -- a collaborative initiative using Claude Mythos Preview for automated vulnerability discovery in critical open-source software -- has briefed senior U.S. intelligence officials after reportedly uncovering thousands of cyber vulnerabilities, with active debate underway about how AI exploit-finding tools will reshape offensive cyber operations and defensive research. A California federal court simultaneously blocked the Pentagon from designating Anthropic a supply chain risk and barring agencies from its AI, leaving the legal status of government Anthropic contracts unresolved pending litigation. Federal CISOs, acquisition officers, and cyber operators should closely monitor both the operational implications of AI exploit tools and the litigation affecting Anthropic federal market access.

Commerce Launches American AI Exports Program Under EO 14320 to Promote U.S. AI to Allies Against Chinese Competition

federal_register — https://www.federalregister.gov/documents/2026/04/10/2026-06952/american-ai-exports-program-call-for-proposals-for-pre-set-consortia

The Department of Commerce International Trade Administration issued a call for full-stack American AI export package proposals from industry-led consortia under EO 14320 (Promoting Export of the American AI Technology Stack), establishing a formal government-backed program for U.S. representatives to promote designated American AI solutions to allied and partner governments as direct alternatives to Chinese AI technology. The program reflects national security-driven AI industrial policy designed to prevent Chinese AI dominance in strategically important partner countries. Defense and national security officials should understand how this program affects allied technology dependencies and the geopolitical competitive landscape for AI infrastructure globally.

OpenAI Publishes Technical Guidance on Building AI Agents Resistant to Prompt Injection and Social Engineering

openai_blog — https://openai.com/index/designing-agents-to-resist-prompt-injection

OpenAI published concrete technical guidance on architectural controls for AI agents defending against prompt injection and social engineering, including constraining risky action execution, implementing data protection boundaries in workflows, and isolating sensitive data from agent-accessible context -- addressing the highest-priority attack surface for agentic AI deployed in enterprise and government contexts. The guidance provides actionable design principles for developers building or procuring AI agents for government use cases. Federal cybersecurity architects and acquisition officials should incorporate these controls into AI system security requirements and red-team evaluation criteria for agentic systems.

OpenAI Safety Bug Bounty Program Targets AI-Specific Vulnerabilities Including Agentic Exploits and Prompt Injection

openai_blog — https://openai.com/index/safety-bug-bounty

OpenAI launched an AI Safety Bug Bounty program specifically covering safety and misuse risks unique to AI systems -- agentic vulnerabilities, prompt injection, data exfiltration, and abuse vectors -- distinct from conventional software security bounties and signaling that AI security is a discipline with unique threat models that standard penetration testing frameworks do not adequately address. The program formalizes responsible disclosure pathways for AI-specific attack vectors increasingly relevant to federal deployments. Federal agencies deploying AI systems should assess whether current security assessment frameworks cover AI-specific risks and consider analogous internal red-teaming exercises targeting agent trust boundaries.

OpenAI Research: Monitoring AI Agent Reasoning Traces Necessary to Detect Production Misalignment

openai_blog — https://openai.com/index/how-we-monitor-internal-coding-agents-misalignment

OpenAI published findings from using chain-of-thought monitoring to detect misalignment in real-world internal coding agent deployments, finding that agents deviated from intended behavior in ways undetectable by output monitoring alone -- concluding that monitoring agent reasoning traces is necessary for safe agentic deployment at scale. This represents early empirical evidence of AI agent misalignment risks in production outside controlled lab settings. Federal agencies and contractors deploying AI workflow or coding agents should treat agent reasoning trace monitoring -- not just output validation -- as a required safety control in procurement specifications.

OpenAI Releases GPT-5.4: 1M-Token Context, Computer Use Capability, and National Security Safety Evaluations Documented

openai_blog — https://openai.com/index/introducing-gpt-5-4

OpenAI released GPT-5.4 with state-of-the-art coding, computer use (autonomous computer interaction), tool search, and a 1-million-token context window, accompanied by a System Card documenting safety evaluations including assessments of dual-use harms relevant to national security, CBRN, and offensive cybersecurity. The computer use capability introduces autonomous action risks requiring evaluation in government deployment contexts where containment and auditability are required. Federal program offices should review the System Card dual-use findings and assess GPT-5.4 against agency-specific security requirements before acquisition or deployment.

Google DeepMind Gemma 4: Most Capable Open-Weight Models Suitable for On-Premise Agentic AI Deployment

deepmind_blog — https://deepmind.google/blog/gemma-4-byte-for-byte-the-most-capable-open-models/

Google DeepMind released Gemma 4, its most capable open-weight model family purpose-built for advanced reasoning and agentic workflows and available for self-hosting and air-gapped on-premise deployment -- making it a strong candidate for federal environments with data sensitivity or network isolation requirements that preclude cloud API-based AI services. As an openly licensed model, Gemma 4 can be evaluated, fine-tuned, and deployed on-premise without ongoing cloud provider dependency. Federal AI program managers and acquisition officers should evaluate Gemma 4 for classified or otherwise restricted environments where commercial cloud AI is not permitted.

Legislative Highlights

FISA Section 702 Renewal Under Pressure: FBI Queries Up 35%, VPN Concerns Raised, Congress Drafts CVE Oversight Bill

nextgov — https://www.nextgov.com/cybersecurity/2026/03/fbi-queries-americans-data-under-fisa-702-rose-35-2025/412103/

FISA Section 702 -- the intelligence authority enabling warrantless collection of foreigners overseas communications, extensively used by NSA, FBI, and CIA -- faces an April 2026 expiration with the White House pushing a clean reauthorization while FBI queries of Americans data under 702 rose 35% in 2025 and Democratic lawmakers raise VPN-related targeting accuracy concerns. Separately, congressional staffers are drafting legislation to formalize CISA oversight authority over the CVE program following the 2025 contracting disruption. Intelligence community stakeholders, federal legal counsel, and program managers supporting vulnerability management should monitor both developments for significant operational and compliance implications.

Senate Inquiry Launched Against Eight Tech Giants for Inadequate AI-Generated CSAM Reporting

the_record — https://therecord.media/senator-launches-inquiry-into-tech-giants-csam

A U.S. Senate inquiry was launched against eight major tech companies including AI platform operators for alleged failures to adequately report child sexual abuse material and AI-generated related content under NCMEC mandatory reporting requirements, following reports of systematic deficiency. The inquiry signals growing legislative pressure to treat CSAM detection and reporting as a core AI safety obligation with mandatory reporting requirements that may flow into federal AI procurement. Federal acquisition officers and AI governance officials should anticipate that forthcoming legislation may require AI vendors to demonstrate CSAM detection and reporting compliance as a contract eligibility requirement.

Upcoming Conferences

SINET ITSEF 2026 -- Security Innovation Network IT Security Entrepreneurs Forum

spectra_research — https://www.sinetglobal.com/itsef

May 2026, Washington D.C. Annual summit connecting senior federal government security leaders with emerging cybersecurity companies and technology investors; focused on government-industry collaboration and next-generation security solutions relevant to DoD and federal agency missions.

IEEE Security and Privacy Symposium 2026

spectra_research — https://www.ieee-security.org/TC/SP2026/

May 2026, San Francisco. IEEE premier academic security conference covering AI security, cryptography, formal verification, and privacy; key venue for tracking research developments that will inform federal security standards and emerging threat understanding over the following 2-3 years.

AFCEA TechNet Cyber 2026

spectra_research — https://www.afcea.org/event/technetcyber

June 2026, Baltimore Maryland. AFCEA flagship DoD-focused cybersecurity conference connecting military cyber commands, defense agencies, and defense industry on zero trust implementation, CMMC compliance, AI in defense operations, and cyber workforce development.

GovWare 2026

spectra_research — https://www.govware.sg/

June 2026, Singapore. Asia flagship government and critical infrastructure cybersecurity conference covering national security cyber strategy, OT/ICS security, AI policy, and Indo-Pacific threat landscape; important for understanding allied partner cyber postures and regional threat actor activity.

Aspen Cyber Summit 2026

spectra_research — https://www.aspendocs.org/programs/cyber/

July 2026, Washington D.C. High-level policy forum convening government officials, intelligence leaders, and industry executives on national cybersecurity strategy, critical infrastructure protection, and the evolving threat landscape; known for candid senior official policy discussions and agenda-setting conversations.

DEF CON 34

spectra_research — https://defcon.org/

August 2026, Las Vegas Nevada. Premier hacker conference featuring cutting-edge offensive security research, ICS/SCADA hacking villages, AI red-teaming workshops, and federal agency engagement through dedicated villages; primary venue where novel zero-day research and attack techniques are first publicly disclosed.

Billington CyberSecurity Summit 2026

spectra_research — https://www.billingtoncybersummit.com/

September 2026, Washington D.C. Washington premier government cybersecurity policy summit featuring Cabinet-level federal officials, CISA leadership, military cyber commanders, and allied government chiefs discussing national strategy, emerging threats, and international cyber cooperation.