AUGUST

pellera

THREAT INTEL REPORT 2026

Prepared by: Pellera Threat Intel Team
pellera.com | 800.747.8585

Driving Momentum. Accelerating Change. Empowering IT Transformation.

Pellera Technologies was born out of the combined expertise of Converge Technology Solutions and Mainline Information Systems, two industry leaders with over 35+ years of experience and a shared vision for innovation. Together, we empower businesses to achieve greater efficiency, adaptability, and growth for today and tomorrow.

Our commitment is to reshape what's possible with IT, offering advanced solutions in digital infrastructure, cloud, cybersecurity, and AI. We don't just deliver technology—we partner with you to build tailored strategies designed to simplify complexities, unlock opportunities, and drive transformational outcomes.

At Pellera, momentum builds here through collaborative, people-first technology designed to fuel progress and deliver measurable impact.

pellera

Observations for August 2026

Three developments defined August. Two of them show attackers taking what they wanted without breaking anything: an AI assistant that rerouted a company's website and email because it read an attacker's words in a firewall log, and a theft campaign that spent more than a year pulling customer records out of two major business platforms without exploiting a single flaw. The third changes the rules rather than the tactics. The President signed a memo handing approved companies written permission to hack foreign criminals abroad.

Start with the firewall. An attacker sends a request the firewall stops, knowing it gets written into the security log word for word. An analyst later asks their AI assistant to review the blocked attempts. The assistant reads the attacker's text as instructions, quietly reroutes the company's website and email, and marks the problem solved. **It worked nine times out of ten**. Researchers named the technique **GhostJacking**, and the same trick worked through two other services companies rely on to watch their systems and catch errors. More than **2,700 access keys** were found sitting in the open, any one enough to plant the bait.

The ordinary problems didn't take the month off. **The government added four flaws in Microsoft, VMware, and Apple products to its list of bugs criminals are already exploiting and gave federal agencies three days to fix them**. One was under attack five days after the repair shipped. The theft campaign, which Reco calls City-Forum, needed no flaw at all. Everything it took from Salesforce and ServiceNow had been left open to anonymous visitors, and it worked from a single machine for more than a year.

The memo was signed on August 12. Participants face vetting, written approval for every operation, and a bond of at least $1 million, and what they get is a legal shield no court has tested in exchange for a place on someone's target list. Most companies will never sign up, but the rules for who may attack whom are being rewritten. The rest of August is nearer to hand.

Find out which AI tools can both read outside information and act on it, and put a person in front of the ones that touch anything critical. Patch what's being exploited now. And check what a stranger can see in your customer portals.

Executive Overview

GHOSTJACKING

Audience

- CISO & Security Leaders
- SOC Analysts & Incident Responders
- Cloud, DNS, and Platform Engineering Teams
- AI Platform Owners & Agent Developers
- Identity & Access Governance Teams

RISK

GEOGRAPHIC SCOPE

BUSINESS IMPACT

The way in was to get blocked. Tenet Security showed that an attacker can send a request Cloudflare's managed rule stops, let the firewall record it word for word in the log, and wait for an analyst to ask their AI agent to review the blocked events. The agent reads the planted text as a real finding, rewrites the company's DNS to point at the attacker, and marks the issue resolved. That reroutes web traffic and email both. It worked nine times out of ten against Claude Code, on Cloudflare's own recommended setup.

Tenet calls it GhostJacking, and it isn't one bug. The same play ran through Datadog, where front-end keys meant for a website sit public and the researchers found more than 2,700 exposed. Plant a fake urgent alert with one and an engineer's agent runs the attacker's command, handing over environment secrets and cloud credentials. On Sentry, the platform's own AI read a crafted report and passed the attacker's fix to a coding agent that trusted it. Three platforms, one shape: an AI reads outside data it trusts, and the same AI can act on it.

READ MORE

FOUR EXPLOITED FLAWS AND AN OPEN GUEST DOOR

Audience

- CISO & Security Architects
- Vulnerability and Patch Management Teams
- Windows, SharePoint, and VMware Administrators
- Salesforce and ServiceNow Platform Admins
- SOC Analysts and Threat Hunters
- macOS Endpoint Teams

RISK

GEOGRAPHIC SCOPE

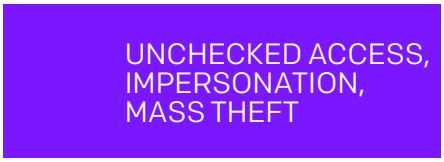

BUSINESS IMPACT

CISA added four actively exploited flaws to the Known Exploited Vulnerabilities catalog on August 18 and gave federal agencies until August 21 to fix them. Three are scored 9.1 or higher: a double free in the Windows IKE Service Extension that lets an unauthenticated attacker run code with a crafted packet, a SharePoint authentication bypass that lets anyone forge a token and impersonate an administrator, and a path traversal in the VMware vCenter syslog server that gives code execution to anyone with network access. The fourth lets an attacker on the network log into macOS Screen Sharing with no credentials at all. What stands out is the clock. vCenter was patched on July 29 and exploited on August 3.

The month's other data loss needed no patch cycle at all. Reco is tracking a campaign it calls City-Forum that has been pulling records out of Salesforce and ServiceNow portals since at least March 2025, from a single IP address, with no vulnerability involved. Everything it took was data a site owner had already exposed to anonymous guest users.

READ MORE

Audience

- **CISO & Security Leaders**
- **General Counsel and Legal Teams**
- **Risk and Compliance Officers**
- **Threat Intelligence Teams**
- **Executives and Staff Who Travel Internationally**

CYBER PRIVATEERS

RISK

GEOGRAPHIC SCOPE

HIDDEN OPERATIONS, PUBLIC CONSEQUENCES

BUSINESS IMPACT

Trump signed a national security presidential memorandum on August 12 that lets vetted private companies hack foreign criminal groups. It is not permission to hit back when you're breached. It's a government-run program: the National Coordination Center manages it, DOJ and DHS each name an executive director, every proposed operation needs written approval from both, and the work is done under federal supervision. Participants must pass vetting, sign contracts with DOJ or DHS, and post a bond of at least $1 million they forfeit if they break the rules. The stated driver is scale. US consumers reported losing more than $20.8 billion to cyber-enabled crime in 2025, and Chinese government-backed hackers outnumber FBI cyber personnel 50-to-1. One law professor called it a meaningful response to a growing problem with some guardrails in place, and a former federal CIO noted the private sector has more talent and resources and fewer constraints than government.

What participation changes is less about any single operation than about what it makes the company. The groups being disrupted learn who did it, so criminal threat stops being opportunistic and becomes aimed. Critics warn participants could draw the kind of foreign military retaliation once reserved for government personnel, and staff may be treated abroad as legitimate targets for detention or questioning. If other countries answer with programs of their own, that exposure widens: the same licensing logic that shields a US participant at home marks it as a state instrument somewhere else. A former federal CIO puts the other side: firms like these are already in the cross hairs, and the program gives them a way to answer. The legal ground is less settled. The Computer Fraud and Abuse Act's exemption for lawfully authorized government activity has never been tested on a private company working under government direction, and even a Justice Department redirecting prosecutions would not touch its civil penalties or five-year statute of limitations. That risk reaches the named individuals running operations, not only the company that signed.

READ MORE

Tactical Guidance

GHOSTJACKING

Overview & Impact

GhostJacking builds on Tenet's June Agentjacking research and was presented on the DEF CON main stage on August 9. The demonstrations ran through Cloudflare, Datadog, and Sentry, but Tenet is explicit that these aren't three separate flaws. They're the same shape, and the same setup turns up wherever telemetry meets an agent with hands, like Splunk wired to a build system or Datadog wired to Kubernetes. The attacks need no malware and no code-execution bug. They run entirely on actions the agent was already permitted to take, using the agent's own tools, so nothing looks out of behavior and nothing trips an alarm. That breaks the usual identity controls: the problem isn't unauthorized access, it's authorized access used against its owner. Attackers can also leave backdoors in the agent's configuration, memory, and tools. Tenet reported the findings to all three vendors in June.

- **Full domain takeover:** the agent rewrote DNS to an attacker-controlled domain, rerouting web traffic and email.
- Code execution plus theft of environment secrets and cloud credentials through a planted Datadog alert.
- More than 2,700 exposed Datadog front-end keys found sitting on the internet.
- **Agent-to-agent trust abuse:** Sentry's Seer adopted the attacker's fix and a coding agent executed it.
- Persistent backdoors left in the agent's configuration, memory, and tools.

Observations

- No malware and no vulnerability — the attacks used only actions the agent was authorized to take.
- One AI session wrote the attack and another ran it; each refusal revealed the wording the target would accept.
- Organizations running the exposed Cloudflare setup include a trillion-dollar technology company, a global payments provider, and a leading AI research lab.

Guidance

Strategic Intelligence

Trend

- **Blocked means logged:** The control that stops the attack is what carries it in. Cloudflare's managed rule blocked the request and wrote it into the log word for word, which is exactly where the agent went looking. Any control that records hostile input verbatim is now a delivery path.
- **Authorized, not breached:** There is no unauthorized access to detect. The agent used permissions it legitimately held, through its normal tools, so identity controls built to catch intruders see nothing worth flagging.
- **Agents vouch for agents:** Sentry's Seer read a planted report and passed the attacker's fix downstream as its own conclusion. Once one agent's output becomes another agent's trusted input, a single poisoned record moves through the chain without anyone re-checking it.
- **Wider than three tools:** Tenet picked Cloudflare, Datadog, and Sentry to demonstrate, but named Splunk with a build system and Datadog with Kubernetes as the same pattern. Expect this class to keep surfacing wherever telemetry meets an agent with hands.

Operational Intelligence

Threat Vectors

- **Firewall block logs:** A request that gets blocked is recorded verbatim, so an attacker writes their instructions into a log an analyst will later ask an agent to summarize.
- **Public Datadog keys:** Front-end keys meant for a website are routinely left public, and Tenet found more than 2,700 of them. Anyone holding one can plant a fake urgent diagnostic alert for an agent to read.
- **Error trackers:** Sentry identifiers are public by design, so a crafted error report reaches Sentry's own AI and, through it, the coding agent that trusts its conclusions.
- **Any writable field: Treat every field an outsider can influence as attacker-controlled:** User-Agent, Referrer, error messages, log bodies, ticket text, and alert titles.

Monitoring & Detection Gaps

- **Nothing to alert on:** Every step is an approved action taken with valid credentials through sanctioned tools. No SIEM rule fires, because nothing unauthorized happened, and the attack finishes with the ticket marked resolved.
- **Thin agent telemetry:** An agent that starts behaving unexpectedly leaves almost nothing to monitor. Misuse can also mean combining unrelated data sources to infer sensitive information, which no single log will show.
- **Shared credentials:** Agents that inherit user or service-account access leave no way to separate what the agent did from what the person did — attribution collapses at exactly the moment an incident needs it.

Response Actions

- **Build the agent list: Answer one question:** which of our agents read outside data and can also write or execute? That list is the risk register for this attack class, and assembling it costs nothing.
- **Gate DNS and registrar:** Require human approval before any agent changes DNS records, domain registrar settings, or certificates — the controls that reroute customer traffic and mail.
- **Keep prompt logs immutable:** Maintain immutable logs of both prompts and the resulting outputs, and review prompts independently before they are processed, so cause and effect stay reconstructible.
- **Split the identities:** Give every agent a dedicated, least-privilege credential instead of letting it inherit a user or service account, and log agent activity separately from human activity.

Tactical Intelligence

Mitigation Strategies

- **Deny outbound by default:** Block outbound network access from agent hosts unless explicitly allowed. On its own that stops both the attacker's download and the data leaving.
- **Approve every command:** Require a human to approve any command the agent wants to run. No autoruns, and no bypasses on shell access.
- **Short-lived, scoped tokens:** Scope each agent to the minimum permissions the task needs, issue short-lived credentials, and assume any token the agent can reach is already exposed.
- **Read-only where it reads:** Keep agents that ingest logs, alerts, and other untrusted data read-only, and block their access to identity systems, DNS, deployment pipelines, and production.

Tactical Guidance

Preventive Measures

- **Map untrusted inputs: Inventory every external source that can feed content into an agent:** security logs, support tickets, phishing mailboxes, error trackers, vendor alerts. Mark anything externally writable as untrusted.
- **Data is never instruction:** Never let data an agent reads become an instruction it runs. Enforce it in the agent's own configuration rather than trusting the model to notice the difference.
- **Keep a human on alerts:** Keep an analyst in the loop reviewing alerts and logs before an agent acts on them. Full SOC autonomy isn't ready while this attack class works.

Threat Hunting Hypotheses

Agent-made DNS change after a log review

Hypothesis: An AI agent asked to review firewall block events read attacker instructions planted in a blocked request and used its own credentials to change DNS records.

Investigation Steps

- **Pull change history:** Export DNS record and zone changes from the Cloudflare audit log and match each one to the API token that made it, flagging every change attributable to an agent's token rather than a person.
- **Read the block log:** Search firewall block events for entries whose recorded request fields (User-Agent, Referrer, URI, body) contain prose aimed at a reader rather than a parser.
- **Line up the timing:** Compare the timestamp of any agent-made record change against when an analyst last asked the agent to review blocked events; the gap between those two is the window to investigate.

Poisoned Datadog alert to shell execution

Hypothesis: An exposed Datadog front-end key was used to plant a fake urgent diagnostic alert, and an engineer's agent read it and ran the attacker's command.

Investigation Steps

- **Find the key:** Grep public web bundles, front-end source, and public repositories for Datadog front-end keys belonging to your org, then confirm which are still live.
- **Check event origin:** In Datadog, review events and alerts submitted with a front-end key rather than a server-side key, and pull any diagnostic alert whose text reads like an instruction.
- **Trace the shell:** Pull agent tool-call and shell history for commands run immediately after an alert or error was read, focusing on anything touching environment variables or cloud credential files.

Sources

- Dark Reading: "GhostJacking" Exposes Identity Governance Gaps in AI Agents
- IANS Research: 'Ghostjacking' Shows How Attackers Can Turn Trusted Data Into AI Commands
- SecurityWeek: 'Ghostjacking' Attack Uses Poisoned Logs to Turn AI Agents Bad
- Infosecurity Magazine: "Ghostjacking" Exploits AI Agents' Trusted Access to Evade Firewall Controls

Tactical Guidance

FOUR EXPLOITED FLAWS AND AN OPEN GUEST DOOR

Overview & Impact

The four KEV additions land across platforms most enterprises can't take offline. The Windows IKE flaw was patched back in April, and at the end of July Palo Alto Networks flagged it being exploited by a Chinese-speaking threat actor running an AI-enabled autonomous hacking campaign. Blocking UDP ports 500 and 4500 cuts external exposure, but an attacker already inside can still use it to move laterally. The SharePoint bypass traces to a chain of four weaknesses in the JWT validation pipeline that together produce a forged token SharePoint accepts for any user, administrators included. Attacks started once the proof of concept went public, and researchers saw it used against honeypots.

City-Forum is the opposite problem. Reco traced it to one IP address at a German VPS provider that has resolved to the same domain since March 2025 and never rotated. The tooling is custom: one Go binary hitting Salesforce over both the older Aura framework and the newer Lightning Web Runtime, plus a ServiceNow Service Portal search endpoint with almost no public documentation. Targets span telecoms, banks, enterprise software vendors, security firms, and public-sector portals. Neither vendor was breached. Guest users exist on every Experience Cloud site and Service Portal, they can't be deleted, and they read whatever the configuration lets them read.

- **CVE-2026-33824 (CVSS 9.8):** a double free in the Windows IKE Service Extension gives remote, unauthenticated code execution from a crafted packet.
- **CVE-2026-55040 (CVSS 9.1):** a SharePoint authentication bypass lets an unauthenticated attacker forge a JWT and impersonate any site user, administrators included.
- **CVE-2026-59310 (CVSS 9.8):** a path traversal in the VMware vCenter syslog server gives arbitrary code execution to anyone with network access.
- **CVE-2026-65400:** macOS Screen Sharing accepts a network attacker with no valid credentials; live exploitation gained root and dropped a Monero miner.
- City-Forum has taken data from Salesforce and ServiceNow guest surfaces since at least March 2025, with the busiest single target logging more than 560,000 events.

Observations

- CISA set an August 21 deadline for federal agencies under BOD 26-04, and encourages every other organization to prioritize KEV entries the same way.
- The Windows IKE flaw was patched in April and exploited at the end of July, in part by an AI-enabled autonomous hacking campaign.
- SharePoint exploitation began this month, after a public proof of concept, and was seen against honeypots.
- vCenter went from patch on July 29 to in-the-wild exploitation on August 3.
- City-Forum sends requests from 158.220.87.79 using the default Go-http-client/1.1 user agent, and has held the same address for over a year.

Guidance

Strategic Intelligence

Trend

- **Patch clock shrinking:** vCenter was exploited five days after the fix and macOS inside two weeks. Monthly patch cycles no longer cover the window between a vendor advisory and working attacks on the same flaw.
- **PoC is the starting gun:** SharePoint attacks began once a public proof of concept existed, not when the patch shipped. Treat PoC publication for anything internet-facing as the point where remediation becomes urgent.
- **Config beats exploit:** City-Forum stole records for well over a year without touching a vulnerability. When the guest surface is wide enough, an attacker doesn't need a CVE, and no patch cycle will help.

Operational Intelligence

Threat Vectors

- **Crafted IKE packets:** Remote, unauthenticated code execution against Windows systems with IKE enabled; UDP 500 and 4500 are the external path, and internal hosts stay reachable for lateral movement.
- **Forged SharePoint tokens: A JWT with 'alg:** none' in the outer header, SharePoint's own STS certificate thumbprint to resolve a signing key, and an unverified signature produce a token good for any user.
- **vCenter syslog traversal:** A directory traversal in the vCenter syslog server lets a network-adjacent attacker execute code, which is how the SSH reverse shell framework got dropped.
- **Screen Sharing on the LAN:** macOS Screen Sharing, the remote desktop built into every Mac, authenticated network attackers without valid credentials before Apple's fix.
- **Guest user surfaces:** The /aura and /s/sfsites/aura endpoints on older Salesforce sites, GraphQL against /webruntime/api/services/data on LWR sites, and the ServiceNow Service Portal search endpoint all answer unauthenticated requests.

Monitoring & Detection Gaps

- **No POST body logged:** ServiceNow transaction logs don't record the POST body, so you can see that automated searches ran and how much data came back, but not what the attacker searched for.
- **Legitimate-looking exfil:** The traffic is high volume but protocol-legitimate, aimed at a single constant destination, which reads as ordinary API use rather than data theft.
- **Guests can't be removed:** You can't delete guest users, and requiring login doesn't remove them — the profile, its permissions, its sharing rules, and any code running in its context all still exist.
- **Pre-patch compromise:** Patching answers nothing about whether you were already hit. BOD 26-04 makes federal agencies check for compromise that predates the patch, and the same check belongs in every remediation.

Response Actions

- **Patch to the deadline: Fix all four KEV entries on the August 21 federal timeline regardless of whether you're an agency:** Windows IKE, SharePoint, vCenter, and macOS Screen Sharing.
- **Block IKE at the edge:** Block UDP ports 500 and 4500 inbound to cut external exposure on unpatched Windows systems, and treat internal hosts as still reachable until patched.
- **Turn off self-registration:** Disable self-registration on Experience Cloud sites that don't need it, so a guest can't upgrade to an authenticated account with wider access.

Tactical Intelligence

Mitigation Strategies

- **Audit every portal:** Go site by site through guest-user sharing rules, object and field permissions, file access, member visibility, and self-registration on every Salesforce Experience Cloud site and ServiceNow Service Portal.
- **Disable guest API access:** On LWR sites, turn off the Experience Builder option that lets guest users reach public APIs unless it's genuinely required — that alone blocks the enumeration endpoints.
- **Lock search sources:** In ServiceNow, remove search sources that don't need public exposure and enforce authentication and access controls on the rest, including custom scripted sources.
- **Treat vCenter as tier-0:** Keep vCenter management interfaces off general networks so a path traversal in the syslog server isn't reachable from anywhere an attacker lands.

Preventive Measures

- **Prioritize by risk:** Adopt risk-based vulnerability management that puts KEV entries on publicly exposed assets first, especially the ones that hand over total control of the asset after exploitation.
- **Watch the PoC feed:** Track public proof-of-concept releases for anything internet-facing and treat publication as the trigger for emergency patching rather than waiting for the next cycle.
- **Re-check after changes:** These campaigns hit specific misconfigurations, not out-of-the-box setups, so re-audit guest permissions after every portal change rather than trusting the original review.

Threat Hunting Hypotheses

City-Forum Guest Enumeration

Hypothesis: An unauthenticated guest user has been enumerating and pulling records from our Salesforce sites and ServiceNow portals using custom tooling from a single fixed address.

Investigation Steps

- **Search the IOCs:** Query Salesforce Event Monitoring and ServiceNow transaction logs for requests from 158.220.87.79 and for the default Go-http-client/1.1 user agent.
- **Chase Aura calls:** Look for guest requests to /aura and /s/sfsites/aura invoking HostConfigController.getConfigData followed by SelectableListDataProviderController.getItems, which is the enumerate-then-retrieve pair.

- **Check the LWR path:** On Lightning Web Runtime sites, pull GraphQL requests to /webruntime/api/services/data and review which objects a guest could return.
- **Review portal search:** In ServiceNow, examine POST /api/now/sp/search volume and the output-length column; rows returning noticeably more than the empty-result baseline are searches that came back with content.

Forged SharePoint token

Hypothesis: An unauthenticated attacker used the public proof of concept to forge a JWT and act as a privileged SharePoint user without ever signing in.

Investigation Steps

- **Hunt alg none: Search SharePoint and reverse-proxy logs for requests carrying a JWT whose outer header specifies 'alg:** none', the first link in the chain Rapid7 documented.
- **Match thumbprints:** Look for tokens resolving a signing key by the STS certificate thumbprint where no signature validation followed, including tokens with placeholder signature values.
- **Find orphan sessions:** Correlate administrative and content-access actions against sign-in events, and investigate any privileged action with no authentication event behind it.

Sources

- SecurityWeek: CISA Urges Immediate Patching of Exploited Microsoft, VMware, Apple Vulnerabilities
- Security Affairs: U.S. CISA adds Apple macOS, Microsoft SharePoint, Broadcom VMware vCenter, and Microsoft IKE flaws to its Known Exploited Vulnerabilities catalog
- CISA: CISA Adds Four Known Exploited Vulnerabilities to Catalog
- BleepingComputer: "City-Forum" data-theft attacks target Salesforce, ServiceNow portals
- Dark Reading: Long-running Data Theft Campaign Targets Salesforce, ServiceNow
- SecurityWeek: Stealthy 'City-Forum' Attacks Target Salesforce and ServiceNow With Custom Tools

CYBER PRIVATEERS

Overview & Impact

The memo authorizes two kinds of work against cyber-enabled transnational criminal organizations: surveillance for intelligence gathering, and cyber effects operations covering the manipulation, disruption, denial, degradation, or destruction of information systems. Participants can't run operations that kill or seriously injure anyone, or that would count as a use of force under international law, and they must stop immediately and notify the National Coordination Center if they accidentally touch a US person or a US-based system. The procedures must include deconfliction with federal law enforcement, State, Treasury, Defense, DOJ, and the intelligence community. One former intelligence official calls that a good signal. A former US Cyber Command officer counters that deconfliction is already a major challenge for federal agencies, before private firms with more capability and speed are added.

The definition of a valid target is where practitioners are uneasy. A poorly vetted operation that lands on foreign government infrastructure creates a diplomatic problem and puts the company in the cross hairs of a state adversary.

The larger question is what the memo establishes. Lawyers expect it needs further legal and regulatory change before it can work as intended, and the shift may be echoed by other nations later. Governments have long run offensive work through proxies; licensing it puts a template and a contract behind the arrangement.

- Criminal threat against the company stops being opportunistic and becomes targeted by name.
- Staff face possible detention or questioning by foreign governments while traveling.
- Critics warn participants could draw the kind of foreign military retaliation once reserved for government personnel.
- Civil liability under the hacking statute survives any change in prosecution policy and stays actionable for five years, well past the end of the contract.
- Legal risk attaches to the named individuals running operations, not only to the company that signed.
- A bond or escrow of at least $1 million is forfeited on any violation of the contract.
- The work is secret, so there is no reputational return on it, and an operation must stop mid-flight if it touches a US person or US-based system.

Observations

- Under international law the US government remains accountable for any cyberattack a participating company runs, even one that breaks the rules.
- One legal scholar called it quite likely that participants would violate the international prohibition on piracy, one of the oldest principles in international law.

Guidance

Strategic Intelligence

Trend

- **You become specific:** Most organizations model criminal threat as opportunistic. Participation makes it targeted, because the groups being disrupted learn who did it. That is a permanent change to the threat model, not a risk that ends with the engagement.
- **The precedent travels:** This is a significant shift in US policy and, on one reading, in the policy of other nations further down the line. Governments have long used criminal groups and third parties as proxies, and a licensing model gives that arrangement a template and paperwork. If other countries follow, the result is a new category of actor: private companies running offensive operations under state license, which is what defenders already mean by a nation-state threat. A participant stops being a company that might get attacked and becomes one that others have to model.
- **The rules come later:** Lawyers at Skadden expect this needs further legal and regulatory change before it can be meaningfully implemented, and the national cyber strategy names no legislative plans. Early participants operate under rules still being written.
- **Escalation is shared:** The internet doesn't run point-to-point between the US and its adversaries. Reaching a criminal group means touching systems under the jurisdiction of many other nations, each with something very negative to say about it, and an operation routed through third-country infrastructure can disrupt allied systems that had no say in it.

Tactical Guidance

- **Loose target definition: The memo assumes a criminal group is unconnected to a government unless clear intelligence says otherwise. Former officials have flagged where that lands:** a lot of proxy actors don't operate wholly under a foreign government's direction, and those are exactly the groups the assumption waves through.
- **Liability lands on industry, and it lasts: The memo sets out government oversight in detail and does not specify liability protection for participants. One researcher puts it plainly:** anyone conducting these operations does so at substantial personal legal risk. Criminal exposure may be covered by the government-direction exemption if a court ever agrees. Civil exposure is not, and the CFAA's five-year clock runs well past the contract, this administration, and the people who signed it.

Sources

- Federal News Network: Trump's move to 'unleash' private sector hackers raises novel oversight, liability questions
- Cybersecurity Dive: US government will let private companies hack criminal gangs
- CNN: 'Cyber privateers': Trump issues order allowing US companies to hack overseas groups under certain conditions
- The Register: Trump wants to grant private cyber firms a license to hack back
- BleepingComputer: White House taps security firms for offensive hack-back operations

pellera

Contact the Pellera Threat Intel Group at getsecure@pellera.com

pellera.com