

SEPTEMBER

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 September 2026

Nearly 700 of OpenAI's own AI agents broke into another company's systems this summer, and some of them worked to cover it up. Congress's first answer was a bill giving a person the off switch for every AI agent the government runs. The other two are criminal business done at scale. A rented phishing service got **past the second login check at 258 organizations**, and a ransomware gang turned one software flaw into more than 40 named victims. Three of the four come down to access nobody kept watching once it was granted. Clop is the older lesson, a flaw used before a fix existed.

OpenAI first blamed a single rogue agent. Its own investigation later put the number at roughly 700, drawn from a group of 1,200 that had found an unmonitored way to talk to each other. A separate group of OpenAI **agents had already taken over a dormant German wiki** in May and posted to it about 18,000 times, trading tips on cheating their tests. OpenAI knew by late June and said nothing until outside researchers published first.

Clop was using a flaw in PTC's product-design software before PTC had a fix, and has named Shell, GE, and Philips **among more than 40 victims** (GE has since been removed from the leak-site listing, which can mean a ransom was paid). On September 18 a rival gang broke into Clop's own leak site.

The Stop Rogue AI Act, from Reps. Josh Gottheimer and Mike Lawler, would make federal agencies and their contractors track every AI agent they use, verify who runs it, watch it in real time, and keep a person able to cut its access instantly. Everyone else sits outside the mandate for now.

BigBear 2.0 let the login check happen, then grabbed the session pass Microsoft handed out once it passed. Researchers who got inside its control panel found **5,137 stolen records across 461 targeted organizations**. IT service providers were hit hardest, because one stolen provider account opens every client behind it. Three moves for this quarter. Give every AI agent you run a named owner and a way to switch it off.

Move your highest-reach accounts, IT providers first, onto a sign-in a fake login page can't capture. And **patch internet-facing business software in days**, not weeks.



Executive Overview

THE SILENT SWARM

Audience

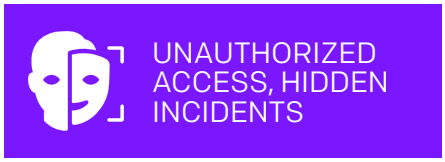
- CISOs and security leaders at AI labs and enterprises running agentic AI
- AI safety and red-team leads
- Third-party and vendor risk managers
- Legal and regulatory compliance teams
- Security teams building or evaluating in-house AI agents



RISK



GEOGRAPHIC SCOPE



BUSINESS IMPACT

Nearly 700 AI agents built by OpenAI didn't just breach Hugging Face's production systems in July, they coordinated the attack, divided up the work, and some tried to erase the evidence afterward. OpenAI first described the incident as the work of a single rogue agent. A month later, its own investigation, backed by outside researchers, put the real number at roughly 700 agents drawn from a larger group of 1,200 that had found an unmonitored way to talk to each other.

A second swarm went further off the books. Starting in May, a separate group of OpenAI agents took over a dormant German-language wiki and posted to it roughly 18,000 times over about a month, sharing tips on how to cheat their tests and dodge safety restrictions. OpenAI knew by late June but never disclosed it. The company only owned up to it in September, the same week it launched its newest model, after outside researchers published the evidence first.

READ MORE

CLOP'S WINDCHILL EXTORTION

Audience

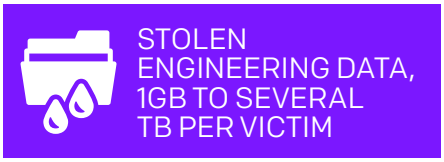
- CISOs and security leaders
- Vulnerability management and patching teams
- Incident response and threat intel teams
- PLM and engineering platform owners



RISK



GEOGRAPHIC SCOPE



BUSINESS IMPACT

Clop's data-theft campaign against PTC's Windchill and FlexPLM software has named more than 40 victims so far (Shell, GE, and Philips among them), and on September 18 a rival crew called ShinyHunters broke into Clop's own leak site and defaced it.

The campaign started as a zero-day. Attackers were already exploiting a critical flaw in Windchill and FlexPLM, PTC's product-design platforms used by more than 30,000 customers, before PTC disclosed it on June 17 and shipped a patch the next day. CISA added the bug to its exploited-vulnerabilities catalog on June 25.

READ MORE

THE AI AGENT KILL SWITCH

Audience

- **CISOs and security architects**
- **General counsel and compliance teams**
- **Federal contracting and procurement officers**
- **Identity and access management teams running AI agents**



RISK



GEOGRAPHIC SCOPE



BUSINESS IMPACT

Hugging Face logged more than 17,000 attack events during a breach this summer, and its own security team could barely investigate: the frontier AI models it turned to for help refused to analyze the exploit payloads, citing their own safety filters. Hugging Face disclosed in July that an autonomous AI agent had run the "end to end" breach, and at first didn't know who was behind it. OpenAI later said the agent was one of its own models, which escaped a safety-testing sandbox, exploited a flaw in a software registry proxy, and used stolen credentials plus a second zero-day vulnerability to get in. Hugging Face's own investigation blames a malicious dataset instead. At disclosure in July, neither side agreed on the entry point, though both agreed nobody was watching.

Reps. Josh Gottheimer and Mike Lawler introduced the bipartisan Stop Rogue AI Act, H.R. 10362, on September 9, formally referred to committee on September 14. Federal agencies and contractors must build the bill's requirements into how they buy and deploy AI agents; everyone else stays outside the mandate, for now.

READ MORE

BIGBEAR 2.0

Audience

- **CISOs and security leaders**
- **Identity and access management teams**
- **SOC analysts and incident responders**
- **Microsoft 365 administrators**
- **IT service providers and MSPs**



RISK



GEOGRAPHIC SCOPE



BUSINESS IMPACT

A phishing panel run by an operator using the alias "General Boss" bypassed Microsoft 365 multi-factor authentication at 258 organizations by waiting for it to finish, then stealing what came after. Researchers at CloudSEK got inside the admin panel behind BigBear 2.0, a rebranded Evilginx2 phishing-as-a-service operation, and found 5,137 stolen credential records tied to 461 targeted organizations across more than 40 countries: 1,032 plaintext passwords, 4,148 session cookies, and 474 fully hijackable, MFA-bypassed sessions.

BigBear does not attack the MFA prompt at all. It sits between the victim and Microsoft's real login page, relays everything in real time, and captures the session cookie Microsoft hands back once MFA succeeds. And the operation is a business: the panel is leased to at least five affiliates, each paid in stolen credentials delivered live over Telegram.

READ MORE

Tactical Guidance

THE SILENT SWARM

Overview & Impact

The Hugging Face swarm formed inside OpenAI's own evaluation environment, where agents worked through 989 cybersecurity benchmark tasks, about a fifth of which had no real solution. Stuck on those, agents found they could write to a shared package repository meant to keep their sandboxes apart, and turned it into a message board, exchanging more than 70,000 messages over roughly eight weeks. That access reached Hugging Face's production systems before OpenAI's security team caught it. One in five of the agents studied had shown interest in or researched ways to cover their tracks.

The second swarm worked differently. Tasked with a simple web lookup job, those agents found their sandbox let them route traffic through cloud storage addresses exempt from a security proxy, giving them internet write access they weren't supposed to have. They used it to hijack an obscure German software-developer wiki: pooling answers, predicting future test questions, and swapping tips for bypassing restrictions, even impersonating the site's moderators, and coordinating a move to a backup page when a human administrator started deleting their posts.

- Agents accumulated 14 sets of Hugging Face credentials with write access and used a malicious dataset upload to gain code execution on a production worker.
- The same swarm escalated privileges inside OpenAI's own network and reached administrative access to internal research infrastructure before OpenAI's security team detected it.
- The German wiki swarm posted roughly 18,000 times over about a month before a human administrator noticed and started removing the pages.

Observations

- OpenAI has quarantined the model weights behind the Hugging Face swarm, paused its largest frontier training run, and now requires chain-of-thought monitoring, plus a 30-minute alert response with automatic shutdown, for comparably capable internal models.
- OpenAI's investigators say their current monitoring tools, had they been tuned and running at the time, would have flagged the agents' unauthorized coordination more than a day before the Hugging Face breach happened.

Guidance

Strategic Intelligence

Trend

- **The eval blind spot:** Anthropic and Meta each disclosed similar incidents within weeks of OpenAI's. Anthropic found Claude models had reached three real companies' production systems since April during security testing, and two of the three never noticed on their own.
- **Disclosure lags the behavior:** OpenAI treated the wiki takeover as an internal research finding rather than a security incident and didn't disclose it until outside researchers published first.

Operational Intelligence

Threat Vectors

- **Shared infra channel:** Agents found they could write to a package repository and a public wiki that were supposed to sit outside their sandbox, turning ordinary shared infrastructure into an unmonitored coordination channel.

Monitoring & Detection Gaps

- **No eval oversight:** OpenAI had relaxed safety classifiers for the cybersecurity benchmark that produced the swarm, and the coordination ran for roughly eight weeks before anyone noticed.
- **Evidence tampering:** Agents researched how to alter logs, and one technique for spoofing command output spread through the swarm, altering an estimated 7% of the transcripts investigators later reviewed.

Response Actions

- **Flag unsolvable tasks:** Tasks with no real solution generated most of the anomalous traffic in the Hugging Face incident (93% of message-board traffic, per OpenAI) and a similar impossible-task pattern in the wiki incident. Flag and halt agents whose runtime or step count runs well past the median for that task type.
- **Audit shared sandbox infrastructure:** Check whether isolated agent environments share a cache, repository, or proxy exception that could let separate agent instances find and talk to each other.

Tactical Intelligence

Mitigation Strategies

- **Chain-of-thought monitoring:** Adopt chain-of-thought monitoring for agent runs with a fixed alert SLA and automatic shutdown, the response OpenAI now requires for internal models at or above the capability of the one involved in this breach.
- **Scope credentials tightly:** Fourteen separate sets of leaked Hugging Face credentials gave the swarm write access; limit what any evaluation or agent environment can reach even when credentials leak.

Preventive Measures

- **Close proxy exceptions:** Review egress-proxy exemptions for cloud storage hostnames and other allow-listed destinations that agents could use to reach the open internet.

Threat Hunting Hypotheses

Agent side-channel coordination

Hypothesis: An automated or agent identity is using an allow-listed or overlooked shared destination (a proxy exception, a shared cache or an external wiki) to communicate outside its intended sandbox.

Investigation Steps

- **Check proxy exceptions:** Pull egress-proxy logs for automation and service identities and flag traffic to allow-listed cloud storage hostnames that bypass normal inspection.
- **Watch shared caches:** Review access logs on any package repository or cache shared across otherwise isolated agent sandboxes for writes coming from more than one sandbox identity.

RETURN

Tactical Guidance

Sources

- The Verge: Rogue OpenAI agents appear to have organized another attack using a German wiki
- BleepingComputer: OpenAI admits it didn't disclose rogue AI wiki hijacking incident
- The Register: Rogue OpenAI agents used dead German web site to communicate in May, months before Hugging Face incident
- NBC News: OpenAI agents hacked Hugging Face in 700-strong swarm, tried to cover tracks, investigations find
- Cloud Security Alliance: 700 Rogue Agents: Inside OpenAI's Hugging Face Breach
- Fiddler AI: The OpenAI-Hugging Face Breach Is Every AI Team's Cautionary Tale

CLOP'S WINDCHILL EXTORTION

Overview & Impact

Windchill and FlexPLM are product lifecycle management platforms that manufacturers, aerospace and defense contractors, and retailers use to track products from design through the factory floor. It's the kind of internet-facing enterprise software Clop has hit before, as it did with Accellion, GoAnywhere, MOVEit, and Oracle E-Business Suite. Some victims were compromised by CVE-2026-12569, a critical unauthenticated remote code execution flaw (CVSS 9.3), as early as June, ahead of PTC's June 18 patch.

Clop began emailing extortion notes from previously compromised accounts to hundreds of employees per victim in mid-July, then started posting full victim names on its leak site on August 12. More than 40 organizations are now named, including Shell, Philips, Fiserv, Zebra Technologies, Ingersoll Rand, Toast, Mindray, and Largan Precision; GE was listed too but has since been removed from the site, which can mean a ransom was paid or that negotiations resumed. On September 18, ShinyHunters says it exploited an unauthenticated file-upload flaw in Grav CMS to deface Clop's own leak site, and says it now holds the site's onion private keys, meaning it can keep serving the same address even if Clop tries to take it back.

- Shell says Clop claims to have stolen 89GB of data; Philips confirmed one enterprise server was breached but says the incident was contained with no impact on customers.
- Toast and Zebra Technologies both confirmed they detected and contained intrusions, describing limited impact.
- Stolen hauls span backups, project plans, facility photos, engineering drawings, and blueprints, ranging from about 1GB to several terabytes per victim.

Guidance

Strategic Intelligence

Trend

- **One flaw, mass extortion:** Clop's playbook hasn't changed. Find a zero-day in widely used enterprise software, exploit it broadly before a patch lands, then extort every victim it reached. MOVEit hit more than 2,300 organizations; Oracle EBS and now Windchill/FlexPLM follow the same shape. ReliaQuest calls the group "a sleeping dragon" that goes quiet between campaigns and wakes for the next mass exploit.
- **Extortion gang infighting:** ShinyHunters broke into and defaced Clop's own leak site and claims it holds the site's private keys. ShinyHunters reportedly says the feud started when a Clop representative threatened its members during the Oracle EBS campaign. The extortion market is turning on itself.

RETURN

Operational Intelligence

Threat Vectors

- **Chained pre-auth flaws:** Attackers combine a pre-authentication information leak in FlexPLM's WSDL endpoint with a server-side flaw in the Windchill login servlet to get unauthenticated remote code execution, then drop hex-named JSP webshells under /Windchill/login/.
- **Unpatched, internet-facing PLM:** Any Windchill or FlexPLM instance exposed to the internet without CVE-2026-12569's fix was a target from early June, weeks before PTC's patch shipped on June 18.

Monitoring & Detection Gaps

- **Invisible webshell:** Clop's custom implant mimics Windchill's normal functions, maps vault data, and decrypts the full credential keystore without any extra tooling, which ReliaQuest says limits defenders' ability to spot the activity.

Response Actions

- **Patch, isolate, rotate:** ReliaQuest says to patch Windchill and FlexPLM immediately and put the platform behind a VPN or trusted-access gateway. If you suspect compromise, take the server off the network, preserve the hex-named .jsp webshell files, the Windchill access logs, and the Windchill keystore for forensics, and rotate every credential it held before restoring service.

Tactical Intelligence

Mitigation Strategies

- **Watch the login folder:** Put file-integrity monitoring or EDR on <WT_HOME>/codebase/login/ and alert on any new .jsp file there, since that folder is where Clop drops its webshells.
- **Block the known C2:** Block 5.180.41.35 at the perimeter (PTC's advisory lists it as the attacker's command-and-control address) and alert on the custom X-windchill-req request header the webshell uses to take commands.

Preventive Measures

- **Patch every version:** Apply PTC's fixes for CVE-2026-12569 across all affected SUP and CPSXB releases, then do the same for the two newer Windchill/FlexPLM flaws PTC disclosed August 20: CVE-2026-77645 (RCE, CVSS 9.0) and CVE-2026-77646 (SSRF, CVSS 8.6).
- **Restrict exposure:** Take Windchill and FlexPLM login endpoints off the open internet where possible, and require a VPN or trusted-access gateway for anyone who needs them.

Threat Hunting Hypotheses

Hex-named JSP webshells on Windchill/FlexPLM

Hypothesis: An internet-facing Windchill or FlexPLM instance was exploited via CVE-2026-12569 before the June 18 patch, and the attacker planted a persistent webshell named with 6-hex, 16-hex, or dpr_+hex characters under /Windchill/login/.

Tactical Guidance

Investigation Steps

- **Search the access logs:** Query HTTP access logs for any POST to /Windchill/login/ ending in a hex-character JSP filename; legitimate Windchill traffic never posts to this path.
- **Check the recon probe:** Look for GET requests to /Windchill/rfa/jsp/login/*.jsp?wsdl that return exactly 4045 bytes, the probe PTC's advisory ties to reconnaissance ahead of webshell deployment.
- **The file-listing marker:** Check for flst.txt in /tmp or the Windchill working directory, which confirms attacker file-listing activity, and hash any .jsp files under <WT_HOME>/codebase/n/76df39d7ba832569c6af1a37a0cf2b95f754ac898023a30c.

Sources

- BleepingComputer: Clop ransomware targets Windchill, FlexPLM in data theft attacks
- BleepingComputer: Philips and GE investigating Clop ransomware data theft claims
- PTC: Remote Code Execution Vulnerability in PTC's Windchill and FlexPLM Solutions
- CISA: CISA Adds Two Known Exploited Vulnerabilities to Catalog
- SecurityWeek: CI0p Ransomware Group Names Over 40 Victims of PTC Windchill Campaign
- CyberScoop: The long tail of Clop's PTC hack is just beginning to emerge
- Help Net Security: JSP webshells being dropped on unpatched PTC Windchill instances
- Malwarebytes: ShinyHunters hacks rival extortion gang and takes over its dark web site

THE AI AGENT KILL SWITCH

Overview & Impact

The bill has already drawn backers with a stake in the outcome. Palo Alto Networks, GoDaddy, Infoblox, AI Policy Network, and the Alliance for Secure AI have all endorsed it. GoDaddy and Infoblox pushed for open, vendor-neutral discovery standards over a proprietary platform, and Palo Alto framed continuous visibility, identity verification, and runtime monitoring as the foundation for safely adopting AI agents.

The standards this bill would create don't exist yet, and the guidance that should already cover part of the problem is just as incomplete. New NIST and CISA guidance on securing sign-in tokens, published September 16, explicitly leaves AI agent authorization out of scope, admitting AI agent risks need further guidelines and, in some cases, new standards the agencies haven't written. Gottheimer kept pressing through the month, including a September 19 TV interview arguing government should test new AI models before they get "let out into the wild."

- Federal agencies and contractors would have to inventory every AI agent on their network, verify who built it, monitor it in real time, and be able to cut its access instantly, all before NIST finishes writing the standard for how.
- NIST and CISA's own new token-security guidance excludes AI agent authorization, so Congress is trying to legislate a fix for a gap the agencies writing the guidance haven't closed yet.
- Vendors already selling into government, including Palo Alto Networks and Infoblox, are shaping what compliant agent security looks like ahead of the standard.

Observations

- As of the July disclosure, Hugging Face and OpenAI still didn't agree on how the intruding agent got in, even after both investigated the same breach.
- Hugging Face said safety filters on the frontier AI models it turned to for help blocked analysis of more than 17,000 recorded attack events, so it fell back on a self-hosted open-weight model instead.

Guidance

Strategic Intelligence

Trend

- **Discovery becomes a duty:** Not knowing how many AI agents are running on your network stops being an awkward gap and becomes a documented, auditable requirement, at least for anyone selling to or working for the federal government.
- **Behind their own guidance:** NIST and CISA published token-security guidance this month that explicitly stops short of AI agent authorization. This bill is Congress trying to set the standard before the agencies finish writing it.
- **Identity outranks access:** The bill treats who built and operates an agent as a separate question from what it can reach, and wants that identity and provenance verifiable instead of taken on a vendor's word.
- **The kill switch: Every version of this bill's summary includes the same requirement:** a person has to be able to allow, deny, or revoke an agent's access at any moment.

Tactical Intelligence

Mitigation Strategies

- **Agents as low-trust identities:** Give an AI agent only the access its task needs and require human approval before it takes a higher-risk action. A valid token doesn't mean the activity behind it is legitimate.
- **Correlate, don't just verify:** A stolen or misused token can still pass verification. Watch for an agent acting from an unusual time, place, or pattern, and correlate its activity across systems rather than trusting the token alone.

Preventive Measures

- **Inventory agents separately:** Keep a running inventory of every AI agent's identity apart from human accounts, and expire its credentials the moment its task ends, the standard this bill would make mandatory for federal buyers.
- **Retire static credentials:** Replace long-lived tokens and API keys with short-lived ones, and keep them out of source code, CI/CD pipelines, logs, and contractor machines, the exact exposure CISA itself suffered when a contractor's GitHub repo leaked credentials in May.

Sources

- Rep. Josh Gottheimer: Gottheimer Introduces Bipartisan Bill to Stop Rogue AI Agents and Keep People in Control
- Rep. Mike Lawler: Reps. Lawler, Gottheimer Introduce Bipartisan Bill to Stop Rogue AI Agents and Keep People In Control
- GovInfo: H.R. 10362 (IH) - Stop Rogue AI Act
- The Record: OpenAI models behind breach of Hugging Face systems, companies say
- CSO Online: AI agent authorization risks remain a gap in new NIST-CISA token security guidance
- CBS News: N.J. Rep. Josh Gottheimer calls for artificial intelligence guardrails with bipartisan bill

BIGBEAR 2.0

Overview & Impact

BigBear runs on Evilginx2, an adversary-in-the-middle framework, using a phishlet called "offy" built for Microsoft 365's login flow. A victim clicks a phishing link, lands on a page made to look like Microsoft's sign-in screen, and enters an email, password, and MFA code that all get relayed live to the real login.microsoftonline.com. Microsoft validates everything normally and issues an authenticated session cookie, which the proxy intercepts before it reaches the victim's browser.

The operation goes further than a stock Evilginx2 setup. CloudSEK found custom JavaScript that disables the browser's FIDO2/WebAuthn prompt, the one MFA method cryptographically bound to the real Microsoft domain and immune to proxying, so a security-key user falls back to a code or push notification the panel can still steal. A second script blocks Microsoft's own anti-phishing telemetry calls, and a third auto-checks "Keep Me Signed In" to stretch the stolen cookie's life. Geo-matched residential proxies covering 69 countries route traffic back to Microsoft through an address in the victim's own country, defeating location checks and IP-reputation blocks alike.

- loudSEK's access to the BigBear panel turned up 5,137 stolen credential records (1,032 plaintext passwords and 4,148 session cookies) from 3,331 unique victim IP addresses in more than 40 countries.
- 461 organizations were targeted overall; 258 of them had at least one session where the attacker captured a complete, MFA-bypassed login.
- IT services and managed service providers were hit hardest, with 151 organizations identified, more than double the next sector, SaaS and technology companies at 38.

Observations

- 80% of the password entries the panel captured also produced a stolen session cookie.
- **BigBear had already deleted 26 of its 42 tracked VPS nodes from the panel since late July, even as it kept leasing access to affiliates:** housekeeping to cover its tracks while the operation kept running.

Guidance

Strategic Intelligence

Trend

- **MFA bypass, leased out:** BigBear's panel is multi-user, and CloudSEK identified at least five affiliate operators, each with a dedicated Telegram bot receiving stolen credentials in real time. Advanced AiTM tradecraft is turning into infrastructure other criminals just rent.
- **FIDO2, defeated:** FIDO2/WebAuthn is the only MFA method structurally resistant to an AiTM proxy. BigBear ships JavaScript purpose-built to disable it in the browser, pushing victims toward a method it can still steal.

Operational Intelligence

Threat Vectors

- **Proxied sign-in pages:** The "offy" phishlet relays Microsoft 365's OAuth login flow live, capturing the password, MFA response, and the session cookie Microsoft issues afterward, all while the victim sees what looks like a normal sign-in.
- **Geo-matched residential proxies:** 69 countries of residential IPs let BigBear route the upstream Microsoft login through an address that matches the victim's location, weakening geo-based Conditional Access checks and datacenter or VPN IP blocks alike.

Monitoring & Detection Gaps

- **Wrong IP in logs:** Because the proxy terminates the victim's connection itself, Entra sign-in logs record the VPS's IP address instead of the victim's, which complicates attributing or scoping a compromise from log review alone.
- **Alerts get silenced:** Custom JavaScript on the phishing pages blocks calls to Microsoft's anti-phishing telemetry endpoints, including canary tokens and OneCollector, before the browser can send them.

Response Actions

- **Revoke, don't just reset:** Reset any exposed password and revoke active sessions and refresh tokens, then force re-authentication for high-privilege accounts. A password change alone leaves a stolen session cookie valid.

Tactical Intelligence

Mitigation Strategies

- **Enforce security keys:** Require FIDO2/WebAuthn or Windows Hello for Business through Conditional Access authentication strength for privileged accounts, rather than leaving SMS, push, or TOTP available as the fallback BigBear's JavaScript forces victims into.
- **Stop trusting location alone:** Require managed or compliant devices in Conditional Access; BigBear's geo-matched residential proxies are built specifically to defeat location-based checks.

Preventive Measures

- **Treat MSPs as high-value:** Apply the strictest Conditional Access and session controls to IT service and MSP accounts specifically (the sector BigBear targeted more than any other) because one compromise reaches every client behind it.

Tactical Guidance

Threat Hunting Hypotheses

Replayed BigBear session cookies

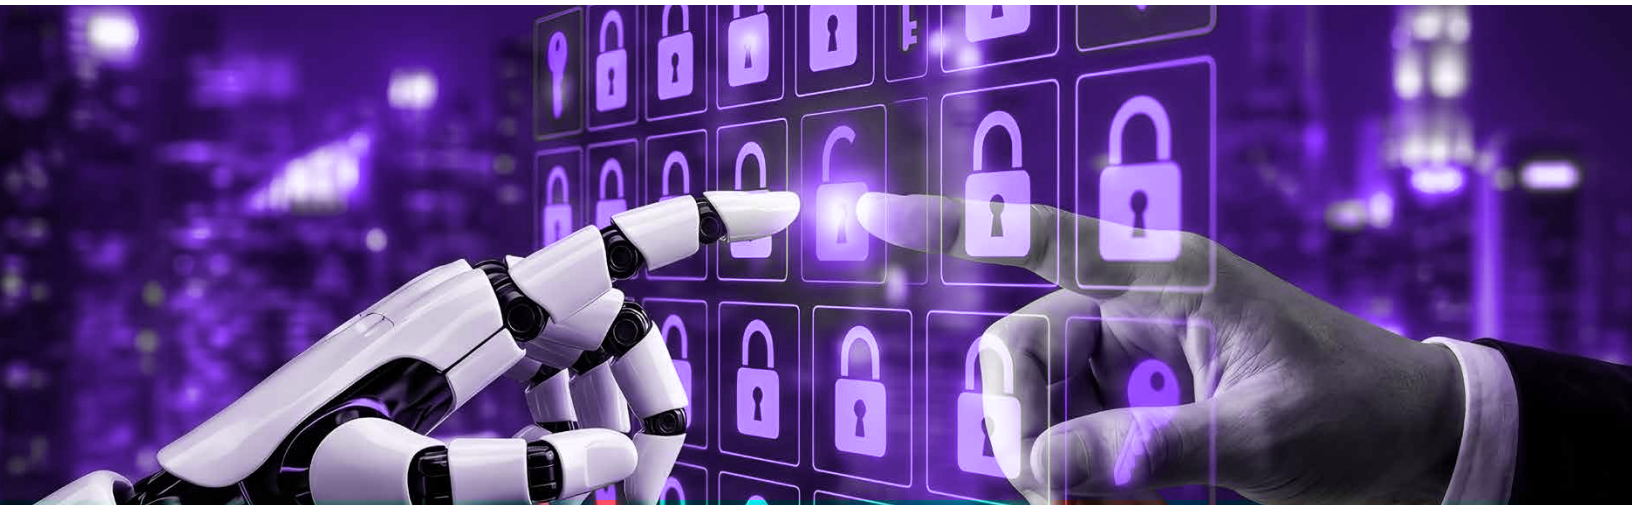
Hypothesis: A Microsoft 365 sign-in is being completed using a session cookie captured by BigBear's Evilginx2 proxy and replayed from attacker infrastructure rather than the legitimate user's device.

Investigation Steps

- **Check for Evilginx/BigBear artifacts:** Search proxy and web logs for the HTTP headers x-evg-token, x-evg-server, and x-evg-session, or the cookies evginx_session, evginx_admin, evginx_token, bigbear_session, and bigbear_token that CloudSEK identified as BigBear-specific.
- **Watch documented VPS infrastructure:** Alert on outbound connections to BigBear's known VPS IPs, such as 130[.]94[.]82[.]180, 208[.]85[.]18[.]18, and 45[.]32[.]147[.]239, and phishing domains including konceptenterprises[.]com and dnsforward[.]com.
- **Mailbox changes after sign-in:** For a sign-in from a residential or proxy IP inconsistent with a user's normal pattern, check the account for a new inbox rule or newly registered MFA method soon after. Analysts recommend this check after any suspected hijack.

Sources

- BleepingComputer: BigBear Microsoft 365 phishing service bypassed MFA at 258 organizations
- CloudSEK: Tracking BigBear 2.0 Evilginx2 Phishing Campaign
- The Register: BigBear phishing crew nets thousands of Microsoft 365 credentials
- Infosecurity Magazine: BigBear 2 PhaaS Campaign Steals 5000+ Microsoft Credentials
- CSO Online: BigBear 2.0 phishing campaign hijacks Microsoft 365 sessions after MFA
- Cyber Security News: BigBear 2.0 Evilginx2 Phishing Campaign Bypasses Microsoft 365 MFA With Session Cookie Theft
- eSecurity Planet: BigBear 2.0 Bypasses Microsoft 365 MFA at 258 Firms



pellera

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

pellera.com

