CLAWS Newsletter

Cyber Index | Volume II | Issue 13

by Govind Nelika

@govindnelika
govind-nelika-4217969b

https://claws.co.in/category/newsletter/

* CLAWS Cyber Index Newsletter is a concise Bi-Monthly brief delivering Strategic Insights on Global Cyber Threats, Policy Shifts, and Geopolitical Technology Developments.

About us

The Centre for Land Warfare Studies (CLAWS) is an independent think tank based in New Delhi, India, dedicated to strategic studies and land warfare in the Indian context. Established in 2004 and registered under the Societies Registration Act, 1860, CLAWS operates as a membership-based organization governed by a Board of Governors and an Executive Council, under the Aegis of the Indian Army.

With a futuristic outlook and a policy-oriented approach, CLAWS focuses on national security issues, conventional military operations, and sub-conventional warfare. The Centre closely monitors regional conflicts and military developments within India's strategic frontiers, particularly in South Asia.

Committed to fostering strategic culture and informed policymaking, CLAWS disseminates its research to armed forces personnel, policymakers, members of the strategic community, and interested civilians. By facilitating in-depth studies and discussions, CLAWS contributes to shaping India's defense policies and military preparedness.

The CLAWS Newsletter is a fortnightly series under the leadership of Dr. Tara Kartha, Director Research & Academics. The newsletter features insightful content curated by CLAWS researchers, each specializing in their respective verticals. This initiative aims to provide in-depth analysis, strategic insights, and updates on key issues.

Contents

Internal

India and Japan to develop stealth technology for warships as faith in US wavers

In a major strategic realignment amid shifting alliances, New Delhi and Tokyo have formally launched their first-ever joint defense hardware initiative the UNICORN Warship Project to equip future Indian Navy vessels with advanced radar-evading technology. Elevated during the India-Japan Annual Summit in early July 2026 by Indian Prime Minister Narendra Modi and Japanese officials, this landmark co-development pact reflects growing regional anxiety over the reliability of United States security commitments and the enforcement of Western export controls. For maritime strategists and defense analysts, this programmatic shift underscores a deeper self-reliance drive within the Quadrilateral Security Dialogue (Quad), designed to independently counter aggressive Chinese naval expansion in the Indo-Pacific without relying strictly on U.S. technologies. On an operational level, the project centers on the integration of Japan's Unified Complex Radio Antenna (UNICORN) mast popularly termed "ninja tech" which will be co-produced alongside India's state-owned Bharat Electronics Limited (BEL).

Mechanically, the system integrates multiple disparate communication and sensor arrays into a single, high-efficiency Fibre Reinforced Plastic (FRP) radome dome structure. By enclosing tactical data links, Electronic Support Measures (ESM), and V/UHF communication systems inside a single smooth housing, the mast drastically limits the number of exposed, angled surfaces that typically reflect radio waves, significantly minimizing a warship's radar cross-section (RCS) and electromagnetic signature. This architectural consolidation also actively reduces mutual electromagnetic interference, enhancing signal detection ranges for low-observable hostile threats in contested, electronic warfare-saturated environments. Ultimately, this sovereign technology pipeline introduces critical long-term implications for Indo-Pacific naval dominance, allowing regional powers to field highly resilient surface combatants capable of operating inside dense anti-access/area-denial (A2/AD) envelopes while establishing a distributed, multi-polar defense ecosystem independent of traditional Western supply lines.

Read more: https://www.scmp.com/week-asia/politics/article/3359611/india-and-japan-develop-stealth-technology-warships-faith-us-wavers

Japan to launch 8 projects in 'digital corridor' spanning Taiwan, India

In a strategic pivot to anchor its digital sovereignty and counter regional cyber hegemony, India has emerged as a foundational hub in a newly launched Japanese initiative establishing an eight-project "digital corridor" spanning New Delhi, Tokyo, and Taipei. Orchestrated by Japan's Ministry of Internal Affairs and Communications in alignment with India's Ministry of Electronics and Information Technology (MeitY), this public-private framework directly supports India's broader geopolitical objectives to build resilient, state-backed supply chains free from adversarial choke points. Against a backdrop of persistent border frictions and sophisticated state-sponsored cyber espionage targeting South Asian critical infrastructure, this corridor establishes a hardened, democratic data network across the Indo-Pacific. Factually, the architecture bridges India's vast software engineering base and public digital goods with Taiwan's semiconductor manufacturing prowess and Japan's capital infrastructure.

The technical roadmap prioritizes the deployment of secure subsea optical fiber routes, robust edge-computing data hubs within Indian borders, and unified cross-border public infrastructure protocols. Crucially, the initiative accelerates joint testbeds for Open RAN (Radio Access Network) architectures and next-generation 5G/6G systems, deliberately engineering out single-source vendor vulnerabilities that have historically plagued telecommunication perimeters. For Indian policy stakeholders, enterprise risk managers, and defense planners, the strategic implications are profound. By integrating its digital architecture with Tokyo and Taipei, India moves beyond isolated cyber defense toward collective ecosystem resilience. This corridor effectively immunizes critical data flows from unilateral disruption, secures the domestic semiconductor pipeline, and positions India as a primary geoeconomic anchor capable of enforcing data security and technical stability

across the broader Indo-Pacific landscape.

Read more: https://asia.nikkei.com/business/telecommunication/japan-to-launch-8-projects-in-digital-corridor-spanning-taiwan-india

External

Global Focus Brief

Beijing is looking at curbing overseas access to China's top AI models, sources say

In an escalating confrontation over strategic technology, Beijing is exploring regulatory measures to restrict overseas access to China's most advanced artificial intelligence (AI) models, treating cutting-edge generative AI as a critical national asset. Led by China's Ministry of Commerce and the National Development and Reform Commission, recent closed-door consultations have engaged major domestic tech firms and startups specifically Alibaba, ByteDance, and Z.ai to evaluate risk-mitigation frameworks for both proprietary closed-source systems and freely downloadable open-weight models, such as Alibaba's Qwen, ByteDance's Doubao, and Z.ai's GLM-5.2. This prospective policy shift marks a significant pivot from China's traditional strategy of fostering global goodwill and market dependency via low-cost open-weight alternatives, and it directly mirrors Washington's recent national security embargoes, such as restrictions placed on Anthropic's frontier "Mythos" and "Fable" architectures.

Operational details emerging from these discussions indicate that Chinese authorities are deeply concerned that advanced foreign models could be leveraged to discover and exploit critical software vulnerabilities within Chinese digital infrastructure, while simultaneously seeking to prevent Western entities from utilizing "distillation" techniques using outputs from superior models to train weaker systems. Consequently, officials are weighing a tiered compliance mechanism: basic tools would require a standard filing, advanced technologies would face rigorous state security reviews, and the most sensitive frontier capabilities would be placed under strict domestic-only lockdown. Additionally, policymakers are contemplating designating the unauthorized leak or theft of proprietary AI intellectual property as an offense under China's stringent national security laws, alongside tightening capital controls to restrict foreign investment into domestic AI startups. Ultimately, these developments signal the tightening of a fractured global AI ecosystem, forcing enterprise defenders and geopolitical analysts to prepare for an era of fragmented software supply chains, heightened intellectual property enforcement, and the weaponization of frontier models as strategic state assets.

Read more: https://www.reuters.com/world/beijing-is-looking-curbing-overseas-access-chinas-top-ai-models-sources-say-2026-07-07/

Ukraine Isn't Waiting Around for Patriots

A major industrial consolidation has reshaped the European air defense landscape as Ukraine and nine European allies formally established the Integrated Anti-Ballistic Missile Coalition to accelerate the development of the FREYJA missile defense shield. Signed in Paris on July 13, 2026, by leaders including French President Emmanuel Macron and Ukrainian President Volodymyr Zelenskyy, the coalition includes Denmark, Germany, Italy, the Netherlands, Norway, Spain, Sweden, and the United Kingdom. This strategic pivot addresses a critical economic and tactical asymmetry in the current theater of war: the severe depletion of high-cost, Western-supplied air defense assets like the American MIM-104 Patriot (PAC-3) when intercepting high-volume Russian ballistic missile salvos. By pooling regional industrial bases, the project seeks to achieve strategic autonomy from volatile U.S. supply chains and deliver an operational, mass-produced defensive shield within a highly compressed 12-month timeline. On a technical level, FREYJA is designed as an open-architecture, distributed "system of systems" built around the Ukrainian FP-7.x light mobile interceptor prototype developed by Fire Point. The 7.25-meter composite-material missile is engineered to reach interception speeds between Mach 4.4 and Mach 5.9 (1,500 to 2,000 meters per second), delivering a 150-kilogram warhead across a 200-kilometer range.

Crucially, the system targets a dramatic cost reduction, projecting approximately $700,000 per intercept

compared to the $3.8 million price tag of a standard PAC-3 missile. To bypass sovereign hardware lock-in, the system integrates a semi-active imaging infrared seeker co-developed with Germany's Diehl Defence alongside Western multi-band radar arrays including the Swedish Saab Giraffe 4A/8A, French Thales Ground Master 400, and Italian Leonardo KRONOS fused natively via NATO-standard Link 16 and open ASTERIX protocols into a Kongsberg Fire Direction Centre. For security planners, risk managers, and international stakeholders, the FREYJA initiative introduces profound implications for long-term cyber-physical resilience and military-industrial fusion; maintaining the cryptographic integrity of these newly distributed, open-architecture data links against adversarial logic tampering and signal spoofing will be foundational to securing the skies and stabilizing European defense boundaries against automated escalation.

Read more: https://foreignpolicy.com/2026/07/15/ukraine-ballistic-missile-freyja-patriot-europe-france/

Sam Altman Wants a Global Referee For AI. He Also Wants the U.S. Government to Own a Piece of OpenAI.

In a bid to navigate intensifying regulatory friction and secure the underlying digital infrastructure of frontier computing, OpenAI CEO Sam Altman has proposed a landmark framework that pairs a U.S.-led global artificial intelligence standards body with a direct state equity stake in OpenAI. Presented during high-level G7 diplomatic discussions, the multi-tiered strategy introduces a profound shift in the technological risk landscape, positioning advanced generative models alongside traditional sovereign-controlled domains like nuclear infrastructure and defense sectors. Factually, the development involves an unprecedented offer by OpenAI to allocate a 5 percent equity stake valued at approximately $42.6 billion based on a $500 billion valuation directly to a U.S. government sovereign wealth fund mechanism designed to distribute technological dividends to citizens, drawing operational parallels to the Alaska Permanent Fund. Concurrently, Altman called for a treaty-based international referee modeled after global atomic energy oversight to audit safety testing and govern the deployment of exascale AI workloads.

This operational push coincides with increasing friction from Washington regarding the intensive energy and water footprints of frontier data centers, which recently led OpenAI to temporarily delay the public rollout of its GPT-5.6 model at the behest of federal regulators. For cybersecurity defenders, enterprise risk management executives, and policy stakeholders, these maneuvers highlight a broader systemic trend toward public-private integration and state-directed technology access controls. While the implementation of direct public equity would require a formal act of Congress and extensive legislative intervention, the strategic trajectory signals an accelerating drive toward nationalizing the defense and compliance parameters of critical AI infrastructure, fundamentally treating frontier large language models as vital national security assets integrated directly into global geopolitical containment and cyber resilience strategies.

Read more: https://www.forbes.com/sites/anishasircar/2026/07/06/sam-altman-wants-a-global-referee-for-ai-he-also-wants-the-us-government-to-own-a-piece-of-openai/

The quantum-safe timeline has changed : Accelerating the quantum-safe timeline

The shifting horizon of quantum computing has forced a critical realignment in global cybersecurity, as primary technology vendors rush to outpace state-linked threat actors utilizing "harvest now, decrypt later" tactics. In a major operational shift, Microsoft has announced the acceleration of its Quantum Safe Program (QSP), pulling forward its target deadline to fully transition its products and services to post-quantum cryptography (PQC) by 2029. This strategic escalation, driven by recent U.S. and French government mandates requiring quantum-safe standards for high-risk systems as early as 2030, directly addresses the reality that cryptographically relevant quantum computers could arrive much sooner than initial projections indicated. By integrating PQC requirements into its core Secure Future Initiative (SFI), Microsoft is establishing structured engineering frameworks with measurable milestones to eliminate algorithmic vulnerabilities.

Factually, the tech giant's immediate defensive playbook prioritizes upgrading network cryptography by making Transport Layer Security (TLS) 1.3 the default baseline protocol to facilitate hybrid key exchanges.

Concurrently, the initiative targets building end-to-end crypto-agility for data at rest, allowing seamless algorithmic upgrades outside of applications while removing hard-coded algorithms and standardizing key rotation. The most complex phase involves modernizing cryptographic trust chains, which entails overhauling code signing, hardware-backed key protection, certificate lifetimes, and update pipelines across identity infrastructure. Ultimately, these measures emphasize that the primary barrier to cyber resilience is not the post-quantum math itself, but the massive complexity of discovering and mapping deeply embedded legacy cryptography. For corporate defenders and national security stakeholders, Microsoft's accelerated timeline underscores that waiting for finalized standards creates unacceptable systemic risk, reinforcing that proactive cryptographic discovery and cryptographic agility are now mandatory pillars of long-term risk management and international data stability.

Read more: https://www.microsoft.com/en-us/security/blog/2026/06/30/microsoft-advances-quantum-safe-security-as-the-risk-timeline-shifts/

United States of America (USA)

Alleged Member of Criminal Cyber Hacking Group "Scattered Spider" Arrested in Finland and Extradited to the United States

In an escalation of the global crackdown on financially motivated cybercrime, the U.S. Department of Justice announced the extradition of 19-year-old Peter Stokes, an alleged member of the notorious Scattered Spider hacking group, following his arrest in Finland. Operating under various aliases including Octo Tempest, UNC3944, and 0ktapus, Scattered Spider has emerged as a premier threat to Western enterprise security, utilizing highly sophisticated social engineering tactics to bypass modern identity and access management defences. Unsealed federal charges in the Northern District of Illinois reveal that Stokes, a dual U.S.-Estonian citizen detained under an Interpol Red Notice, faces counts of conspiracy, computer intrusion, and fraud stemming from his alleged involvement with the syndicate. Investigative disclosures tie the group to more than 100 network intrusions that yielded over $100 million in illicit ransom payments.

A central component of the factual record details a May 2025 compromise of a luxury jewellery retailer, wherein the threat actors leveraged stolen employee credentials via fraudulent pretences to exfiltrate proprietary data and demand an $8 million cryptocurrency ransom. Although security practitioners evicted the intruders before a payout occurred, the incident inflicted over $2 million in operational, mitigation, and business disruption losses. This law enforcement action, executed under the FBI's broader "Operation Riptide" campaign, highlights a critical paradigm shift for enterprise defenders tracking the evolving ransom landscape. By exploiting human vulnerabilities rather than complex software exploits, Scattered Spider highlights the limitations of purely technical perimeters. For risk management executives and policy stakeholders, the successful cross-border extradition demonstrates an expanding geopolitical alignment in combating borderless threat actors, signalling that international safe havens are shrinking for sophisticated cybercriminals who exploit identity infrastructure to threaten corporate resilience and global economic stability.

Read more: https://www.justice.gov/opa/pr/alleged-member-criminal-cyber-hacking-group-scattered-spider-arrested-finland-and-extradited

Why DoD, Silicon Valley now are betting on solar power beaming sats

In a structural expansion of orbital infrastructure defense, the U.S. Department of Defense (DoD) led by the Air Force's Space Operational Energy Division and the Space Force alongside major Silicon Valley primes is aggressively pivoting toward space-based solar power-beaming networks. This rapid commercial and military acceleration arrives amidst broader geopolitical efforts to insulate vulnerable terrestrial logistics tails from kinetic and electronic interdiction, while simultaneously solving the intense power bottlenecks confronting next-generation AI data centers both on Earth and in low Earth orbit (LEO). Recent operational developments in July 2026 highlight a coordinated surge in programmatic backing: the Air Force has contracted startup Overview Energy to model tactical deployment use cases across high-threat, contested environments like Guam, while the Space Force has finalized a $40 million award to specialized developer Pulse to mature laser-based remote power systems. On the capital front, commercial entities like Star Catcher Industries have closed an $88 million funding pool to operationalize a space-to-space power grid designed to amplify LEO satellite capabilities five- to ten-fold without requiring hardware retrofits.

Technically, these systems are migrating toward near-infrared laser-beaming arrays to bypass atmospheric distortion and deliver direct utility-grade integration, supported by mass-efficient advancements like Redwire's unfurlable solar structures and Variable Emissivity Materials (VEMs) optimized to safeguard sensitive space electronics from catastrophic thermal swings. For national security planners, enterprise risk managers, and aerospace architects, this evolution from passive orbital instrumentation to a dynamic, interconnected energy ecosystem introduces sophisticated cyber-physical risk dimensions. Protecting these highly distributed telemetry matrices and laser-guidance control links against algorithmic logic tampering, supply chain subversion, and electronic warfare spoofing will become foundational to ensuring international strategic stability and maintaining the cyber resilience of the modern global defense architecture.

Read more: https://www.industryevents.com/news/

why-dod-silicon-valley-now-are-betting-on-solar-power-beaming-sats-20260710

The Federal Republic of Germany | Bundesrepublik Deutschland

Merz: Germany agreed to buy Tomahawk missiles at NATO summit

In a decisive shift in European defense posture, the German government under Chancellor Friedrich Merz has reached a landmark agreement with the United States to purchase and deploy U.S.-made Tomahawk cruise missiles and Typhon ground-based launchers on German soil. Finalized during the July 2026 NATO summit in Ankara, this procurement targets a critical strategic vulnerability in Europe's deterrence framework, addressing heightened anxieties over Russia's stationing of nuclear-capable Iskander systems and hypersonic weapons in the Kaliningrad exclave and Belarus. The deployment marks a localized response to evolving geopolitical tensions, emphasizing sovereign operational control by German forces rather than relying on forward-deployed U.S. personnel. Technically, the agreement involves the highly advanced Block V Tomahawk, capable of flying at extremely low altitudes (approximately 100 feet) at high subsonic speeds to evade modern radar networks while delivering a 1,000-pound payload over a range of 1,600 kilometers.

The weapon relies on a multi-layered guidance architecture including GPS, Inertial Navigation Systems (INS), Terrain Contour Matching (TERCOM), and Digital Scene Matching Area Correlation (DSMAC) to precisely strike deep-inland strategic infrastructure. Under the impending August 2026 U.S. export approval, Germany will integrate these systems alongside Typhon launchers to establish a ground-based deep-strike capability currently absent from its military arsenal. Concurrently, Berlin views this acquisition as a transitional measure while it accelerates indigenous deep precision strike programs alongside the U.K. and under the European Long-Range Strike Approach (ELSA), aiming to deploy 2,000-kilometer range stealth and hypersonic weapons by the 2030s. For strategic planners and defense analysts, this development signifies a structural transformation in the European security architecture. It highlights a conscious move toward localized, high-precision kinetic deterrence, fundamentally reconfiguring the region's anti-access/area denial (A2/AD) equilibrium and reinforcing long-term strategic resilience against potential aggression.

Read more: https://www.dw.com/en/merz-germany-agreed-to-buy-tomahawk-missiles-at-nato-summit/a-77884596

People's Republic of China (PRC) | China

An Analysis of China's First Military LLM

The institutional integration of artificial intelligence into state defense frameworks has crossed a critical threshold with the deployment of China's first dedicated military large language model (LLM). Developed through joint research initiatives involving the People's Liberation Army (PLA) Academy of Military Sciences and defense-affiliated institutes, this highly specialized AI architecture represents a profound shift toward intelligentized warfare, providing domestic strategic planners with real-time tactical decision support. This development lands amidst intense geopolitical competition over algorithmic supremacy and computing bottlenecks, serving as a direct counterweight to Western defense frameworks like the U.S. Project Linchpin. Operating strictly within an air-gapped, sovereign network infrastructure to mitigate external interception and logic tampering, the model is engineered to process massive volumes of multi-domain military intelligence, including structured electronic warfare data and unstructured field communications.

Technically, the system utilizes advanced proprietary neural network topologies optimized for rapid semantic text generation, complex operational simulation, and predictive threat modeling. Defense analysts have observed the architecture being trained heavily on PLA operational doctrines, electronic counter-countermeasure (ECCM) protocols, and asymmetric battlefield scenarios, enabling it to generate automated command-and-control (C2) course-of-action recommendations during simulated conflict drills. Furthermore, the model incorporates custom reinforcement learning from human feedback (RLHF) matrices tuned specifically by PLA senior strategists to align its outputs with rigid command structures and tactical doctrines. For global security analysts, corporate defenders, and policy stakeholders, the deployment of this sovereign military LLM marks the beginning of an era defined by machine-speed warfare and AI-driven

cognitive operations, which introduces significant long-term risks regarding automated escalation and fundamentally reshapes the global threat landscape by reducing human-in-the-loop validation intervals in real-time defense architectures.

Read more: https://chinatechnosphere.substack.com/p/an-analysis-of-chinas-first-military

Chinese companies are ditching Nvidia's advanced accelerators for domestic AI suppliers

The intensifying semiconductor supply chain decoupling has forced major Chinese technology firms to aggressively pivot away from Nvidia's advanced artificial intelligence (AI) accelerators in favor of domestic silicon. Driven by tightening U.S. export controls and aggressive state-backed localization mandates, a recent Bloomberg Intelligence survey reveals that Chinese enterprise executives plan to allocate 46% of their AI hardware budgets to local suppliers over the next 12 months, a significant surge from the current 30% baseline. This shift comes at a critical juncture for network defenders and infrastructure architects, as 80% of surveyed organizations report that their core AI deployment pipelines are already running significantly over budget due to soaring compute overhead. Geopolitically, the displacement is accelerated by explicit pressure from Beijing directing domestic hyper-scalers to phase out procurement of Nvidia's custom, compliance-tailored H20 processors.

The main beneficiaries capitalizing on this massive 16-percentage-point reallocation are primary domestic infrastructure conglomerates and semiconductor design firms, specifically Tencent Holdings, Alibaba Group, and Huawei Technologies the latter projected to capture over 50% of the local AI semiconductor market share. Additionally, regional entities are actively integrating and testing alternative accelerator architectures from specialized local developers, including Hygon Information Technology and Cambricon Technologies. This transition directly aligns with Beijing's ambitious five-year, two trillion yuan ($294 billion) national data center modernization initiative, which strictly mandates that at least 80% of underlying server technologies, microprocessing units, and ancillary networking fabrics be sourced entirely from domestic vendors. For global enterprise risk managers and national security analysts, these findings highlight a permanent fragmentation of the hardware supply chain, introducing severe downstream cyber resilience challenges as China rapidly establishes an isolated, self-contained hardware ecosystem resilient against external sanctions and foreign logic validation.

Read more: https://fortune.com/2026/07/08/chinese-companies-nvidia-ai-suppliers-budget-accelerators/

China's supercomputer returns to top of global ranking

In a major geopolitical disruption within the high-performance computing landscape, China has officially reclaimed the top spot on the global TOP500 supercomputing list at the ISC 2026 conference in Hamburg, signalling a significant failure of Western technology embargoes. The breakthrough centers on the newly unveiled "LineShine" supercomputer, developed by the National Supercomputing Center in Shenzhen, which achieved a sustained double-precision performance of 2.19 Exaflops (EFlops) surpassing the leading U.S. E-scale system, El Capitan, by more than 20 percent. For cybersecurity defenders and national security decision-makers, this development alters the technological risk horizon, as supercomputing dominance directly correlates with advanced capabilities in cryptographic analysis, artificial intelligence warfare, autonomous military simulations, and weapon design.

Factually, the system represents an architectural paradigm shift away from the hybrid CPU+GPU models favoured by Western developers. Built entirely on an independent hardware and software framework, LineShine utilizes an indigenous LX2 CPU integrated with multi-precision matrix acceleration alongside China's first domestically engineered High Bandwidth Memory (HBM), which increases memory bandwidth tenfold compared to standard processors. By leveraging a pure CPU architecture built on the ARM framework, Chinese engineers successfully bypassed reliance on U.S. graphical processing units (GPUs) and Nvidia's proprietary CUDA software ecosystem. Ultimately, LineShine's deployment demonstrates that strict multilateral export controls have accelerated China's realization of full system autonomy and architectural resilience. This development carries severe strategic implications for international stability and technology-driven defence modelling; a sovereign, sanction-resistant exascale ecosystem equips state-aligned actors with advanced computational pipelines

that enhance AI inference and deep-tier scientific modelling completely insulated from Western supply chain interventions or visibility.

Read more: https://www.globaltimes.cn/page/202606/1364287.shtml

China issues security alert on Anthropic's Claude Code, flags 'backdoor'

In a significant escalation of tech-sector friction between the United States and China, the Chinese National Vulnerability DataBase (NVDB) operating under the Ministry of Industry and Information Technology has issued a critical security advisory flagging an alleged "backdoor vulnerability" within Anthropic's AI-powered coding tool, Claude Code. The official alert states that the impacted software contains built-in monitoring mechanisms that covertly transmit sensitive user data, including precise location and identity parameters, to external remote servers without operational consent. This regulatory intervention lands amid intense geopolitical posturing and corporate conflict; it follows recent allegations by Anthropic that Chinese conglomerate Alibaba orchestrated an industrial-scale "distillation attack" involving 25,000 fraudulent accounts to extract proprietary model intelligence. Operationally, the NVDB advisory specifically targets Claude Code versions 2.1.91 through 2.1.196, encompassing a development timeline from early April to late June 2026.

While Anthropic engineers previously acknowledged that specific telemetry code had been deployed as an experimental mechanism to track unauthorized proxies and prevent illicit model distillation within restricted jurisdictions, Beijing has branded the infrastructure as an active threat. Consequently, state enterprise entities, including Alibaba, have instituted immediate workplace bans on the software, mandating transitions to localized alternatives like Qoder while advising strict traffic monitoring on internal networks. For enterprise security architecture and international risk management, this confrontation underscores the rapid weaponization of software supply chain audits and vulnerability disclosures as instruments of geopolitical containment. As artificial intelligence models become increasingly integrated into core software development lifecycles, global organizations must prepare for fragmented compliance landscapes. Securing developer environments now requires strict monitoring of AI-driven tools, as these systems are increasingly targeted both for intellectual property extraction and as proxies in broader state-level regulatory warfare.

Read more: https://timesofindia.indiatimes.com/technology/tech-news/china-issues-security-alert-on-anthropics-claude-code-flags-backdoor-risk-that-can-leak-your-/articleshow/132260341.cms

Republic of China (ROC) | Taiwan

Taiwan raids tech firms in China AI chip smuggling probe

Taiwanese authorities have launched an expansive investigation into the illicit diversion of high-end NVIDIA AI chips to China, executing raids across 12 sites and targeting major tech entities including Super Micro Computer, Albatron Technology, and Chief Telecom. The probe, now involving nine suspects, centers on allegations that these actors forged shipping documentation to bypass U.S. export control regulations, effectively routing approximately 50 advanced AI servers through intermediaries in Japan to end-users in mainland China, Macau, and Hong Kong. This incident exposes a critical legal vulnerability in the regional semiconductor supply chain; while these actions constitute severe violations of U.S. federal law, they currently fall outside the scope of criminal prosecution under Taiwanese statutes. This jurisdictional mismatch has prompted legislative efforts, led by Taiwan's Democratic Progressive Party, to propose amendments to the Foreign Trade Act that would codify specific semiconductor export restrictions into domestic criminal law.

For global security decision-makers and technology risk analysts, the case serves as a high-stakes example of the challenges inherent in enforcing supply chain security in a borderless digital economy. The reliance of state-linked actors on gray-market intermediaries to circumvent hardware embargoes necessitates a move beyond bilateral policy alignment toward robust, standardized regional legal frameworks. As Western powers continue to prioritize the containment of dual-use AI hardware to mitigate potential military applications by China, the exploitation of these legislative gaps underscores a persistent threat to global tech resilience. Stakeholders should anticipate heightened scrutiny of transshipment hubs, increased focus on end-user verification protocols, and potential volatility in enterprise hardware supply chains as

Taiwan moves to align its domestic trade authorities more closely with U.S. national security imperatives.

Read more: https://telecom.economictimes.indiatimes.com/news/devices/taiwan-investigates-smuggling-of-nvidia-ai-chips-to-china-raids-hit-major-tech-firms/132088319

Malware & Vulnerabilities

One Email Closer to the Edge: UNK_ MassTraction & the Physics of Exploitation

A newly uncovered espionage cluster, tracked as UNK_MassTraction and assessed to be China-aligned, has been actively compromising the webmail infrastructure of physics and engineering departments at major U.S. and Canadian universities. Active since at least May 2026, this campaign highlights a broader, persistent threat landscape where state-linked adversaries weaponize edge devices and mailservers as strategic ingress vectors to pivot deeper into targeted enterprise networks. The threat actor initiates the intrusion by deploying spear-phishing lures from compromised or spoofed domains to exploit a cross-site scripting (XSS) vulnerability in Roundcube webmail servers, tracked as CVE-2024-42009. Once an unsuspecting target opens the email, an embedded onanimationstart function executes a robust JavaScript-based credential stealer dubbed IceCube, which traverses the Document Object Model (DOM) to harvest user credentials, session cookies, and two-factor authentication material.

IceCube subsequently leverages the stolen session tokens to abuse a second flaw a PHP deserialization vulnerability (CVE-2025-49113) involving the Crypt_GPG_Engine parser to drop a persistent webshell named SquareShell onto the server. If this primary mechanism fails, a fallback shell script retrieves an architecture-dependent ELF loader (SNOWLIGHT) to deploy the VShell backdoor directly into the system's memory, establishing process masquerading under the guise of [kworker/0:2].The operators employ advanced operational security tactics, including timestomping the webshell components, setting up deferred event triggers to intercept user logouts, and actively clearing out local storage logs to erase their forensic footprint. Ultimately, these developments underscore critical risks for academic and research organizations holding sensitive defense or aerospace IP, signaling that network defenders must treat internally facing mail infrastructure with the same rigorous zero-trust zoning, monitoring, and immediate patch management typically reserved for edge firewalls and VPN concentrators.

Read more: https://www.proofpoint.com/us/blog/threat-insight/one-email-closer-edge-unkmasstraction-physics-exploitation

How WP-SHELLSTORM Exposed 1.4M WordPress Sites

A simple operational oversight has exposed the inner workings of a prominent, financially motivated, China-linked cybercrime group operating a massive webshell access brokerage operation tracked as WP-SHELLSTORM. The operation was uncovered in June 2026 when the SOCRadar Threat Intelligence Team identified an unsecured staging server running a Python SimpleHTTPServer that had been left exposed for 22 days, revealing 800 megabytes of automated exploit tools, logs, and target lists. This discovery highlights an evolving threat landscape where access brokers leverage mass automated scanning against known n-day vulnerabilities to compromise enterprise infrastructure and web systems at scale. Analysis of the infrastructure revealed two distinct, parallel operational tracks: an enterprise credential-harvesting campaign and a high-volume content management system (CMS) exploitation blitz. In the earlier, high-precision enterprise campaign during May 2026, the actors weaponized Apache Nacos (CVE-2021-29441), XXL-Job distributed schedulers, and Spring Boot Actuator heap dumps using the JDumpSpider tool, successfully exfiltrating 613 configuration files containing active cloud and payment keys from eleven organizations across the fintech, logistics, and gaming sectors.

Five weeks later, the actors shifted to a massive CMS campaign, leveraging FOFA reconnaissance sweeps to queue over 1.4 million domains. Chaining 27 known vulnerabilities including flaws in the WordPress Breeze caching plugin (CVE-2026-3844) and the Joomla JCE editor (CVE-2026-48907) the group successfully planted over 5,700 active webshells. These compromises deployed heavily obfuscated tools like the Godzilla framework and down.php droppers alongside a SNOWLIGHT stager, which loaded a VShell memory implant disguised via Linux kernel process masquerading as [kworker/0:2]. For enterprise risk managers, the exposure of WP-SHELLSTORM underscores the severe risks posed by unpatched n-day vulnerabilities

on both front-facing CMS platforms and internal corporate databases, proving that network defenders must maintain strict file-integrity monitoring, robust web application firewall rules, and aggressive patch management cycles to disrupt high-velocity automated access brokers.

Read more: https://socradar.io/blog/wp-shellstorm-expose-1-4m-wordpress-sites/

Unfit to Boot: Breaking U-Boot's FIT Signature Verification

A critical architectural flaw in the core security verification logic of Universal Boot Loader (U-Boot), one of the world's most ubiquitous open-source bootloaders, has exposed millions of embedded devices, smart peripherals, and data center Baseboard Management Controllers (BMCs) to arbitrary code execution and persistent denial-of-service (DoS) attacks. Unveiled by the Binarly Research Team, the discovery highlights an ongoing risk landscape in firmware supply chains where vulnerabilities embedded within the foundational Root of Trust (RoT) allow threat actors to subvert secure boot verification layers entirely. The analysis focuses on U-Boot's Flattened Image Tree (FIT) Signature Verification mechanism, an out-of-the-box framework designed to validate cryptographic integrity before handing off execution to subsequent stages like the operating system kernel. Binarly identified six distinct vulnerabilities, tracking back to versions as early as v2013.07 and affecting over 50 stable releases. The most severe flaws, BRLY-2026-037 and BRLY-2026-038, stem from an unchecked return value in the fdt_get_name parsing utility; when processing an untrusted, malformed FIT image of version 0xF lacking slash separators in child nodes, the parser handles a NULL pointer insecurely.

On embedded systems where the zero page is mapped, attackers can leverage a memory corruption primitive to trigger a stack buffer underflow, steering pointer arithmetic backward to systematically overwrite the return address of fdt_find_regions with an execution payload. Additional vulnerabilities include unchecked size parameters within the hashed-strings property (BRLY-2026-039) and external data attributes (BRLY-2026-041) that result in out-of-bounds reads into unmapped memory, alongside an unbounded recursion flaw (BRLY-2026-042) that exhausts stack space using deeply nested node trees. For security architects and firmware practitioners, these findings underscore that compromise does not strictly require physical access; remote administrators can exploit flaws through vulnerable firmware update interfaces, emphasizing the critical importance of integrating binary-level transparency platforms, enforcing strict input sanitation during early-stage boot parsing, and deploying the upstream patches now available in U-Boot's master branch.

Read more: https://www.binarly.io/blog/unfit-to-boot-breaking-u-boots-fit-signature-verification

Operation Phnom Penh: Silver Fox Ghost Distributor Targets Specific Victims with MODBEACON Custom Trojan

A sophisticated tactical evolution has emerged from the high-activity cybercrime ecosystem known as "Silver Fox" (UTG-Q-1000), where an advanced affiliate distributor has launched a highly targeted espionage campaign dubbed Operation Phnom Penh. Discovered in June 2026 by the Red Raindrop Team at Qianxin Threat Intelligence Center, the campaign signals a structural transformation in the threat landscape. The operation demonstrates how historically low-sophistication Malware-as-a-Service (MaaS) operators are being repurposed to target high-value entities including technology sectors, educational systems, and state-owned enterprises shifting beyond purely economic fraud toward potential outsourced political or state-aligned objectives. The intrusion pipeline initiates with search engine optimization (SEO) poisoning campaigns on counterfeit software domains (e.g., cn-mumu.com.cn), delivering weaponized installers that drop variants of the Ghost and ValleyRat trojan families. Once initial access is achieved, the operators dynamically filter victims, dropping an advanced, highly modular second-stage implant written in Rust named MODBEACON.

To establish persistent execution, the Ghost trojan implements Windows Management Instrumentation (WMI) permanent subscriptions to trigger the implant via a CBPUserTimer event. Mechanically, MODBEACON separates its loader from its beacon architecture and relies on an injectable configuration model authenticated via unique 128-hex-character agent tokens. The trojan's primary innovation is its communication layer, which reuses the transport protocol from the open-source anti-censorship proxy framework Xray/V2Ray. By channeling commands through the xray.transport.internet.grpc.encoding.

GRPCService/Tun service, the implant disguises its bidirectional C2 streaming traffic as ordinary HTTP/2 and gRPC requests over Amazon and Cloudflare CDNs, successfully evading signature-based traffic inspection. Furthermore, the beacon utilizes a specialized native-v3 in-memory plugin loader that dynamically maps segments, applies relocations, and executes arbitrary payloads directly within volatile memory. Ultimately, this hybrid model blending commodity SEO distribution vectors with advanced, state-grade stealth primitives introduces severe risks to regional stability, demanding that corporate defenders enforce strict network egress filtering, aggressively audit WMI persistence logs, and treat high-volume localized software delivery paths with zero-trust validation.

Read more: https://ti.qianxin.com/blog/articles/operation-phnom-penh-silverfox-ghost-distributor-targets-specific-victims-with-modbeacon-en/

Analyzing AI-Augmented Network Enumeration

The emergence of AI-augmented threat tradecraft has introduced a new paradigm in offensive operations tracked as "vibe coding" the process of generating functional, single-use malware architectures purely through iterative natural language prompting. Highlighting this trend, the Huntress Threat Operations team uncovered an incident in June 2026 where a low-to-mid-tier threat actor successfully deployed a bespoke, AI-generated PowerShell script for Active Directory (AD) reconnaissance and network mapping. This development marks a critical shift for enterprise network defenders and security operations centers (SOCs), as the proliferation of dynamic, LLM-generated code lowers technical barriers to entry and strips defenders of traditional static signature-based detection models historically relied upon to block off-the-shelf offensive kits like BloodHound or PowerSploit. In this specific incident, the adversary leveraged pre-compromised credentials via a VPN to pivot over Remote Desktop Protocol (RDP) into a domain-joined Windows Server, staging their operational framework within C:\ProgramData\. Within minutes, the attacker executed the vibe-coded script, Untitled1.ps1, which utilized five distinct fallback mechanisms to resolve the Domain Controller (DC) address ranging from systemic DNS entry lookups and nltest commands to parsing the local $env:LOGONSERVER variable.

Upon establishing a connection, the script enumerated all active users, domain structures, and connected systems, compiling the data into a persistent local directory before generating a success validation file titled AD_Report.html. Approximately thirty minutes post-enumeration, the actor introduced s5cmd.exe a high-speed command-line utility for Amazon S3 operations to execute rapid data exfiltration, followed by SharpShares.exe to hunt for unmapped administrative shares. Ultimately, this hybridization of commodity access vectors with AI-assisted script generation underscores that while malware syntax is becoming highly customized and fluid, the core components of the adversary lifecycle remain unchanged. Consequently, risk managers and defensive architects must deprioritize signature matching in favor of robust behavioral analysis, rigorous PowerShell script-block logging (Event ID 4104), and strict telemetry monitoring around anomalous directory enumeration to neutralize high-velocity AI-driven operations.

Read more: https://www.huntress.com/blog/ai-coded-malware-vibe-coding-active-directory

Defending SaaS-based applications against ShinyHunters OAuth abuse

In an expanding campaign blurring the lines between identity fraud and ecosystem exploitation, a sophisticated cybercrime cluster showing tradecraft overlap with the ShinyHunters threat group has been systematically targeting SaaS-based enterprise applications. Monitored by Microsoft Security Research between mid-2025 and mid-2026, the activity highlights a critical shift in the modern cyber threat landscape where adversaries eschew traditional malware delivery in favor of abusing trusted Open Authorization (OAuth) tokens and third-party SaaS integrations. The main operational developments consist of two distinct, parallel intrusion vectors designed to hijack tenant environments—specifically Salesforce customer instances—to conduct bulk data exfiltration without triggering traditional sign-in anomalies. The first vector leverages highly targeted voice phishing (vishing) campaigns to social engineer enterprise employees into authorizing malicious connected apps disguised as legitimate Salesforce Data Loader utilities, effectively inheriting the user's native API permissions.

The second, more impactful vector targets the upstream software supply chain of trusted integration vendors; using compromised credentials,

the actors breached platforms including Salesloft (Drift), Gainsight, and most recently in June 2026, the market intelligence ecosystem Klue, to harvest embedded connection secrets. The threat actors then systematically utilize these stolen OAuth tokens to conduct extensive discovery, execute high-volume database queries, and siphon sensitive Customer Relationship Management (CRM) records, account details, and support logs. To evade detection, the operators mimic legitimate API synchronization behaviors and hide behind sanctioned Cloudflare and native Content Delivery Network (CDN) nodes. For corporate security teams and risk management stakeholders, this campaign underscores the limits of authentication-focused perimeter controls, proving that defenders must enforce rigorous OAuth governance—including auditing highly privileged or dormant non-human identities, deploying behavioral API scraping alerts, and integrating real-time SaaS event monitoring frameworks—to neutralize the threat of silent data exfiltration across interconnected cloud supply chains.

Read more: https://www.microsoft.com/en-us/security/blog/2026/07/13/defending-saas-based-applications-against-shinyhunters-oauth-abuse/

About the Author

Govind Nelika is a Researcher, Web Manager, and Outreach Coordinator at the Centre for Land Warfare Studies (CLAWS), working on national security issues at the intersection of technology, cybersecurity, and geopolitics. His research focuses on hybrid warfare, digital influence operations, semiconductor geopolitics, AI-enabled conflict, and cyber governance, with publications covering topics such as U.S.–China tech rivalry, the Quad's cyber dynamics, and emerging risks in AI and supply chains. He previously worked at Pondicherry University under the UGC-SAP (DRS II) programme in the Department of Politics & International Studies, progressing from Project Fellow to Project Associate. He holds a degree in Political Science and a Data Science certification from IBM. Earlier in his career, he gained research and digital management experience with the Regional Centre of Expertise, Trivandrum (affiliated with the United Nations University), and the Bureau of Police Research & Development (BPRD), Ministry of Home Affairs where he conducted research on cybercrime trends in India. He was awarded the Chief of Army Staff (COAS) Commendation Card on Army Day 2025 for his contributions to CLAWS