Top Tier Target | What It Takes to Defend a Cybersecurity Company from Today's Adversaries

sentinelone.com/labs/top-tier-target-what-it-takes-to-defend-a-cybersecurity-company-from-todays-adversaries/

Tom Hegel

Executive Summary

- In recent months, SentinelOne has observed and defended against a spectrum of attacks from financially motivated crimeware to tailored campaigns by advanced nation-state actors.
- These incidents were real intrusion attempts against a U.S.-based cybersecurity company by adversaries, but incidents such as these are neither new nor unique to SentinelOne.
- Recent adversaries have included:
  - DPRK IT workers posing as job applicants
  - ransomware operators probing for ways to access/abuse our platform
  - Chinese state-sponsored actors targeting organizations aligned with our business and customer base
- This report highlights a rarely-discussed but crucially important attack surface: security vendors themselves.

Overview

At SentinelOne, defending against real-world threats isn't just part of the job, it's the reality of operating as a cybersecurity company in today's landscape. We don't just study attacks, we experience them firsthand, levied against us. Our teams face the same threats we help others prepare for, and that proximity to the front lines shapes how we think, and how we operate. Real-world attacks against our own environment serve as constant pressure tests, reinforcing what works, revealing what doesn't, and driving continuous improvement across our products and operations. When you're a high-value target for some of the most capable and persistent adversaries out there, nothing less will do.

Talking about being targeted is uncomfortable for any organization. For cybersecurity vendors, it's practically taboo. But the truth is security vendors sit at an interesting cross-section of access, responsibility, and attacker ire that makes us prime targets for a variety of threat actors, and the stakes couldn't be higher. When adversaries compromise a security company, they don't just breach a single environment—they potentially gain insight into how thousands of environments and millions of endpoints are protected.

In the past several months alone, we've observed and defended against a spectrum of attacks ranging from financially motivated crimeware to tailored campaigns by advanced nation-state actors. They were real intrusion attempts targeting a U.S.-based cybersecurity company — launched by adversaries actively looking for an advantage, access, or leverage. Adversaries included DPRK IT workers posing as job applicants, ransomware operators probing for ways to access/abuse our platform, and Chinese state-sponsored actors targeting organizations aligned with our business and customer base.

We are certainly not the only ones facing these threats. In the spirit of furthering collective defenses and encouraging further collaboration, we're pulling back the curtain to share some of what we've seen, why it matters, and what it tells us about the evolving threat landscape—not just for us, but for every company building and relying on modern security technology.

DPRK IT Workers Seeking Inside Jobs

One of the more prolific and persistent adversary campaigns we've tracked in recent years involves widespread campaigns by DPRK-affiliated IT Workers attempting to secure remote employment within Western tech companies– including SentinelOne. Early reports drew attention to these efforts and our own analysis revealed further logistical infrastructure to launder illicit funds via Chinese intermediary organizations. However, neither gave a sense of the staggering volume of ongoing infiltration attempts. This vector far outpaces any other insider threat vector we monitor.

These actors are not just applying blindly — they are refining their process, leveraging stolen or fabricated personas, and adapting their outreach tactics to mirror legitimate job seekers in increasingly convincing ways. Our team has tracked roughly 360 fake personas and over 1,000 job applications linked to DPRK IT worker operations applying for roles at SentinelOne — even including brazen attempts to secure positions on the SentinelLabs intelligence engineering team itself.

Public reporting of DPRK IT workers applying to threat intelligence positions

Engagement and Adversary Interaction

Instead of staying passive, we made a deliberate choice towards intelligence-driven engagement. In coordination with our talent acquisition teams, we developed workflows to identify and interact with suspected DPRK applicants during the early phases of their outreach. This collaboration was key. By embedding lightweight vetting signals and

monitoring directly into recruiting processes — without overburdening hiring teams — we were able to surface anomalous patterns tied to DPRK-affiliated personas piped directly into our Vertex Synapse intelligence platform for analyst review.

Our attempted interactions offered rare insights into the craftiness and persistence of these infiltration campaigns — particularly the ways in which adversaries adapt to the friction they encounter.


[redacted] Referral Request for Full Stack Engineer at SentinelOne [redacted]

[redacted]@gmail.com>

to me

Hi Tom

I hope you're doing well. I came across the Full Stack Engineer position at SentinelOne and wanted to reach out. I'm [redacted] a Senior Full Stack Developer with 10+ years of experience in Java, Spring Boot, Angular, AWS, and DevOps.

Would you be open to referring me? If you'd like, we can set up a quick call so I can share more about my experience before you decide. Let me know if that works for you—I'd really appreciate your help!

Best,

[redacted]

https://www.linkedin.com/in/[redacted]

Inbound DPRK referral request to strategic employees

The attackers are honing their craft beyond the job application and recruitment process. An operation of this scale and nature requires a different kind of backend infrastructure, such as a sprawling network of front companies to enable further laundering and logistics.

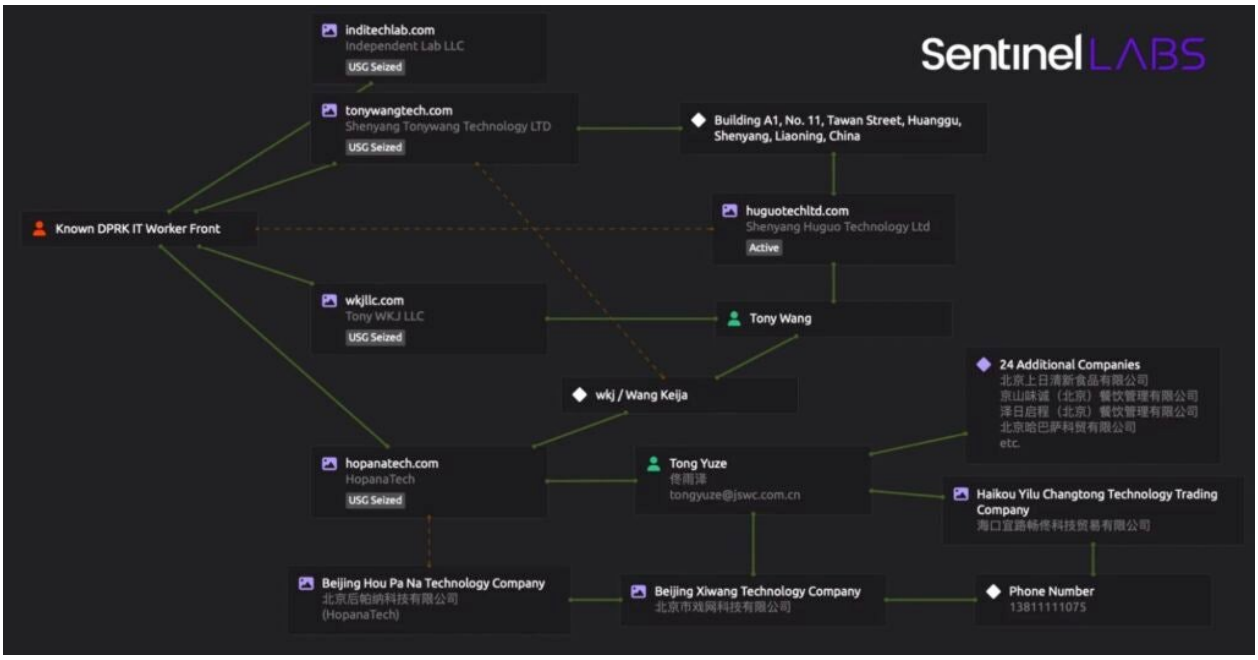

DPRK IT Worker Front Company Network (November 2024)

Helping Hiring Teams Help Us

A key takeaway in working on this investigation was the value of intentionally creating inroads and sharing threat context with different teams not normally keyed into investigations. Rather than cluelessness, we encountered an intuitive understanding of the situation as recruiters had already been filtering out and reporting ‘fake applicants’ within their own processes.

We brought campaign-level understanding that was combined with tactical insights from our talent team. The payoff was immediate. Recruiters began spotting patterns on their own, driving an increase in early-stage escalation of suspicious profiles. They became an active partner that continues to flag new sightings from the frontlines. In turn, we are codifying these insights into automated systems that flag, filter, enrich, and proactively block these campaigns to lower the burden on our recruiters and hiring managers, and reduce the risk of infiltration.

Make cross-functional collaboration standard operating procedure: equip frontline business units—from recruiting to sales—with shared threat context and clear escalation paths so they can surface anomalies early without slowing the business. Codifying insights with automation will consistently bring bi-directional benefits.

The DPRK IT worker threat is a uniquely complex challenge — one where meaningful progress depends on collaboration between the security research community and public sector partners.

Ransomware Group Capability Development

Financially motivated threat actors frequently target enterprise security platforms —products designed to keep them from making money—for direct access. SentinelOne, like our peers, is no exception. While uncomfortable, this is a reality the industry faces continually and should handle with both transparency and urgency.

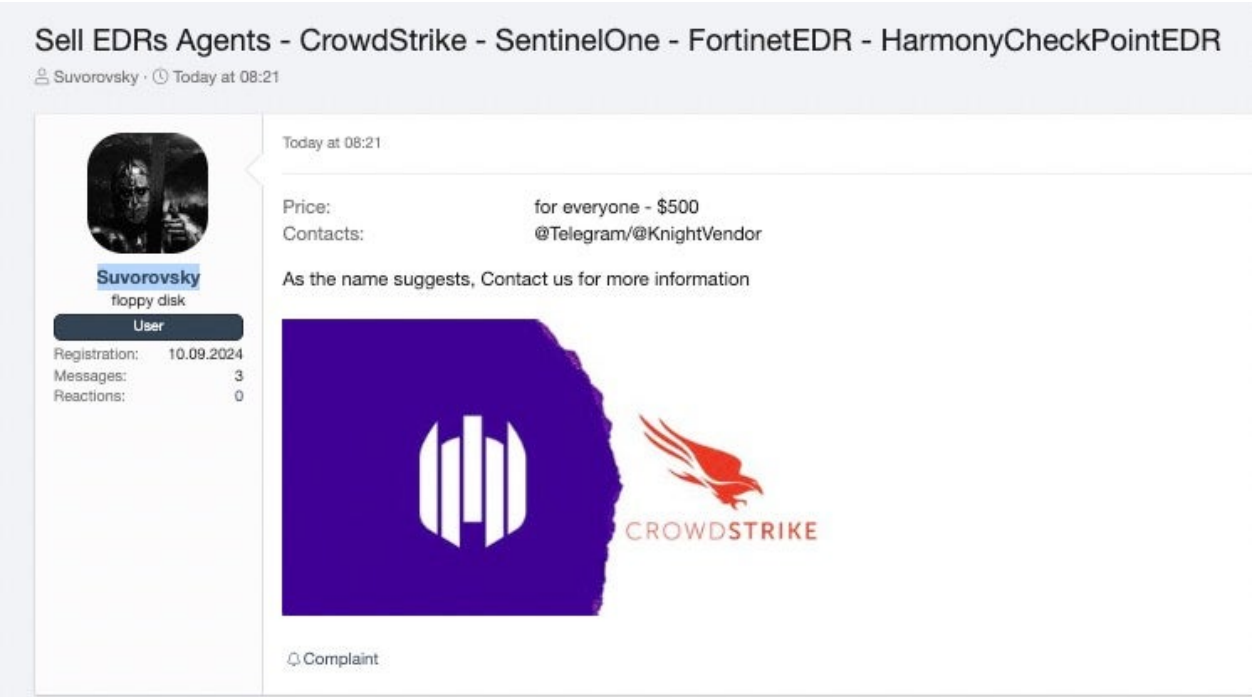

Forum post offering security product access

Privileged access to administrative interfaces or agent installers for endpoint security products provides tangible advantages for adversaries seeking to advance their operations. Console access can be used to disable protections, manipulate configurations, or suppress detections. Direct, unmonitored access to the endpoint agent offers opportunities to test malware efficacy, explore bypass or tampering techniques, and suppress forensic visibility critical for investigations. In the wrong hands, these capabilities represent a significant threat to both the integrity of security products and the environments they protect.

This isn't a new tactic. Various high-profile criminal groups have long specialized in social engineering campaigns to gain access to core security tools and infrastructure—ranging from EDR platforms (including SentinelOne and Microsoft Defender) to IAM and VPN providers such as Okta. Their goal: expand footholds, disable defenses, and obstruct detection long enough to profit.

Recent leaks related to Black Basta further underscore this trend. The group's operators were observed testing across multiple endpoint security platforms—including SentinelOne, CrowdStrike, Carbon Black, and Palo Alto Networks—before launching attacks, suggesting a systematic effort to evaluate and evade security tools prior to deployment.

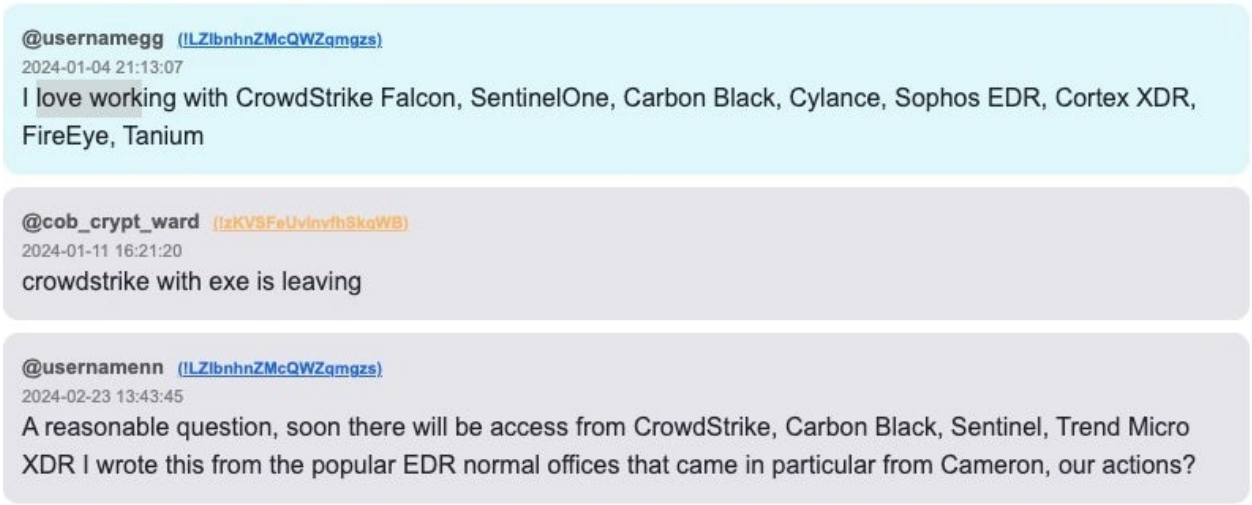

Black Basta leak excerpts

Economy/Ecosystem

There is an increasingly mature and active underground economy built around the buying, selling, and renting of access to enterprise security tools. For the right price, aspiring threat actors continually attempt to obtain time-bound or persistent access to our EDR platform and administrative consoles. Well-known cybercrime forums are filled with vendors openly advertising such access—and just as many buyers actively seeking it. This includes long-established forums like `XSS[.]is`, `Exploit[.]in` and RAMP.

That said, more of this activity has been moving to confidential messaging platforms as well (Telegram, Discord, Signal). For example, Telegram bots are used to automate trading this access, and Signal is often used by threat actors to discuss nuance, targeting and initial access operations.

This supply-and-demand dynamic is not only robust but also accelerating. Entire service offerings have emerged around this ecosystem, including “EDR Testing-as-a-Service,” where actors can discreetly evaluate malware against various endpoint protection platforms.

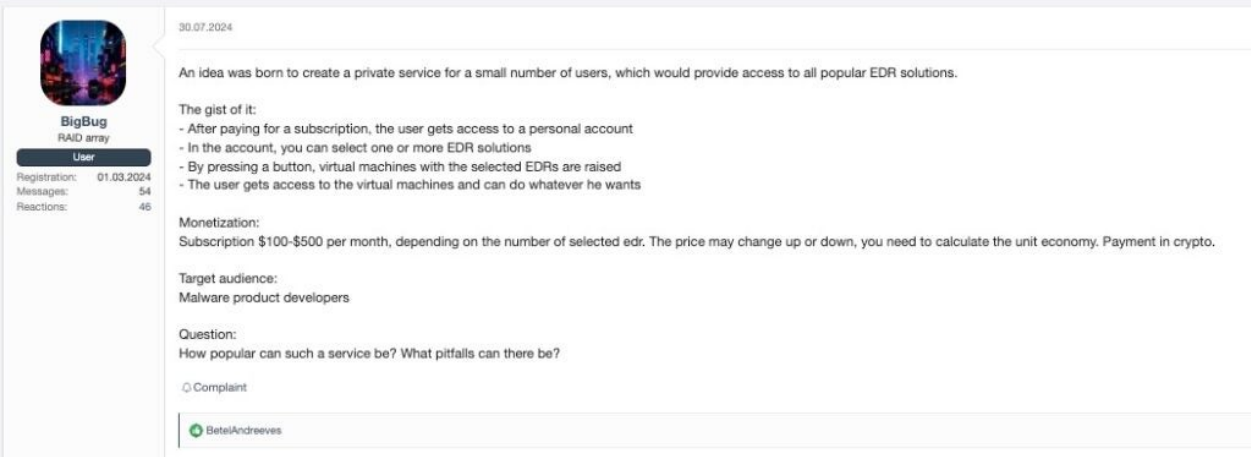

Proposed Private EDR testing service

While these testing services may not grant direct access to full-featured EDR consoles or agents, they do provide attackers with semi-private environments to fine-tune malicious payloads without the threat of exposure—dramatically improving the odds of success in real-world attacks.

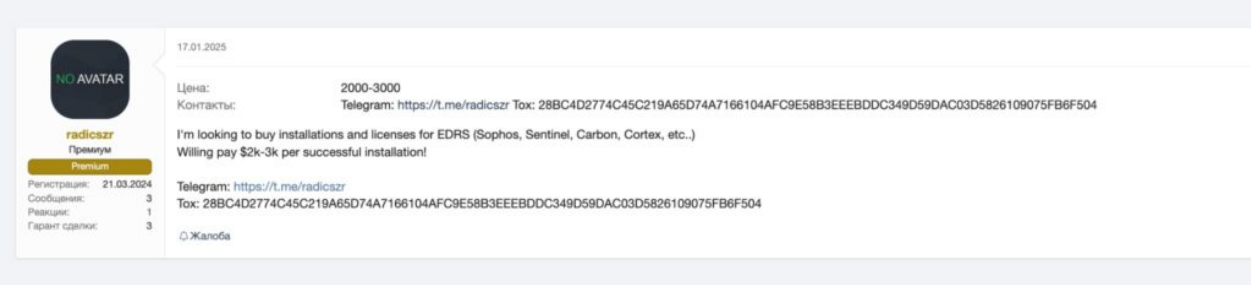

Prospective buyer for EDR installs

Access isn’t always bought, however. Threat actors frequently harvest legitimate credentials from infostealer logs—a common and low-cost method of acquiring privileged access to enterprise environments. In cases where existing customers reuse credentials, this can translate into a threat actor also gaining access to security tools. In more targeted operations, actors have also turned to bribery, offering significant sums to employees willing to sell out their account access.

These insider threats are not hypothetical. For instance, some groups have been observed offering upwards of $20,000 to employees at targeted companies in exchange for insider assistance—an approach openly discussed in the same dark web forums where compromised credentials and access are routinely traded.

On the defensive side, this requires constant monitoring and maintenance. Situational awareness has to be prioritized in order to maintain platform integrity and protect our legitimate customers. Our research teams are constantly monitoring for this style of abuse and access 'leakage', focusing on anomalous console access and site-token usage, and taking necessary actions to revoke these access vectors. This prohibits threat actors from fully interacting with the wider platform, and essentially orphans leaked agent installs, limiting the use of the agent in the hands of the threat actor.

Nitrogen — Threat Operators 'Leveling Up'

Some ransomware operations are now bypassing the underground market altogether—opting instead for more tailored, concentrated-effort impersonation campaigns to gain access to security tools. This approach is epitomized by the Nitrogen ransomware group.

Nitrogen is believed to be operated by a well-funded Russian national with ties to earlier groups like Maze and Snatch. Rather than purchasing illicit access, Nitrogen impersonates real companies—spinning up lookalike domains, spoofed email addresses, and cloned infrastructure to convincingly pose as legitimate businesses. Nitrogen then purchases official licenses for EDR and other security products under these false pretenses.

This kind of social engineering is executed with precision. Nitrogen typically targets small, lightly vetted resellers—keeping interactions minimal and relying on resellers' inconsistent KYC (Know Your Customer) practices to slip through the cracks.

These impersonation tactics introduce a new layer of complexity for defenders. If a threat actor successfully acquires legitimate licenses from a real vendor, they can weaponize the product to test, evade, and potentially disable protections—without ever having to engage with criminal markets.

This highlights a growing challenge for the security industry: reseller diligence and KYC enforcement are clearly part of the threat surface. When those controls are weak or absent, adversaries like Nitrogen gain powerful new ways to elevate their campaigns—often at a lower cost and lower risk than the black market.

Lessons Learned and Internal Collaboration

One of the most impactful lessons from tracking adversaries targeting our platform has been the value of deep, early collaboration across internal teams — particularly those not traditionally pulled into threat response efforts. For example, by proactively engaging with our reseller operations and customer success teams, we can surface valuable signals on questionable license requests, reseller behavior anomalies, and business inconsistencies that could have otherwise gone unnoticed.

By creating shared playbooks, embedding lightweight threat context, and establishing clear escalation paths, reactive processes turn into proactive signal sources. Now, suspicious licensing activity—especially when paired with evasive behaviors or mismatched domain metadata—can surface much earlier in the workflow.

To scale this effort, we increasingly lean into automation. By codifying threat patterns—such as domain registration heuristics, behavioral metadata mismatches, and reseller inconsistencies—organizations can automate enrichment and risk-scoring for incoming licensing requests. This can then be used to dynamically filter, flag, and in some cases, auto-block high-risk activity before it reaches onboarding.

The growing trend of adversaries exploiting sales processes—whether through impersonation, social engineering, or brute-force credential use—means security vendors must treat every access vector, including commercial and operational pipelines, as part of the attack surface. Making cross-functional threat awareness standard operating procedure and integrating detection logic at the edge of business systems is essential.

We're continuing to improve this work in quiet ways. And while we won't share every detection logic here (for obvious reasons), we encourage others in the industry to pursue similar internal partnerships. Sales and support teams may already be seeing signs of abuse —security teams just need to give them the lens to recognize it.

Chinese State-Sponsored Adversaries

One notable set of activity, occurring over the previous months, involved reconnaissance attempts against SentinelOne's infrastructure and specific high value organizations we defend. We first became aware of this threat cluster during a 2024 intrusion conducted against an organization previously providing hardware logistics services for SentinelOne employees. We refer to this cluster of activity as PurpleHaze, with technical overlaps to multiple publicly reported Chinese APTs.

The PurpleHaze Activity Cluster

Over the course of months, SentinelLABS observed the threat actor conduct many intrusions, including into a South Asian government supporting entity, providing IT solutions and infrastructure across multiple sectors. This activity involved extensive infrastructure, some of which we associate with an operational relay box (ORB) network, and a Windows backdoor that we track as GoReShell. The backdoor is implemented in the Go programming language and uses functionalities from the open-source reverse_ssh tool to establish reverse SSH connections to attacker-controlled endpoints.

SentinelLABS collectively tracks these activities under the PurpleHaze moniker. We assess with high confidence that PurpleHaze is a China-nexus actor, loosely linking it to APT15 (also known as Nylon Typhoon, or other various outdated aliases). This adversary is known for its global targeting of critical infrastructure sectors, such as telecommunications, information technology, and government organizations – victimology that aligns with our multiple encounters with PurpleHaze.

We track the ORB network infrastructure observed in the attack against the South Asian government organization as being operated from China and actively used by several suspected Chinese cyberespionage actors, including APT15. The use of ORB networks is a growing trend among these threat groups, since they can be rapidly expanded to create a dynamic and evolving infrastructure that makes tracking cyberespionage operations and their attribution challenging. Additionally, GoReShell malware and its variations, including the deployment mechanism on compromised machines and obfuscation techniques have been exclusively observed in intrusions that we attribute with high confidence to China-nexus actors.

ShadowPad Intrusions

In June 2024, approximately four months prior to PurpleHaze targeting SentinelOne, SentinelLABS observed threat actor activity targeting the same South Asian government entity that was also targeted in October 2024. Among the retrieved artifacts, we identified samples of ShadowPad, a modular backdoor platform used by multiple suspected China-

nexus threat actors to conduct cyberespionage. Recent ShadowPad activity has also included the deployment of ransomware, though the motive remains unclear — whether for financial gain or as a means of distraction, misattribution, or removal of evidence.

The ShadowPad samples we retrieved were obfuscated using ScatterBrain, an evolution of the ScatterBee obfuscation mechanism. Our industry partner, Google Threat Intelligence Group (GTIG), have also observed the use of ScatterBrain-obfuscated ShadowPad samples since 2022 and attribute them to clusters associated with the suspected Chinese APT actor, APT41.

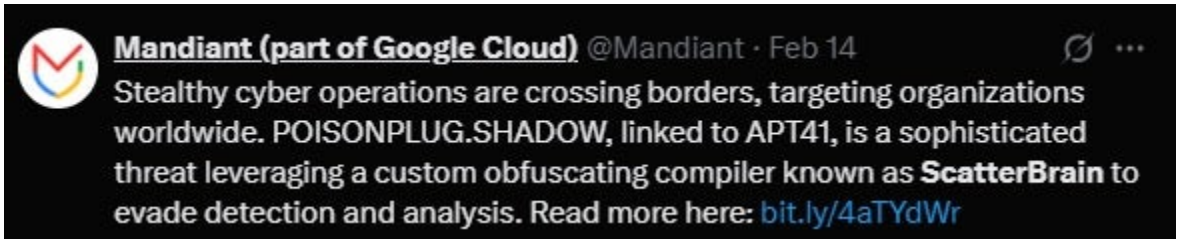

GTIG APT41 Use of ScatterBrain

Investigations continue in determining the specific actor overlap between June 2024 ShadowPad intrusions and the later PurpleHaze activity. We do not rule out the involvement of the same threat cluster, particularly given the extensive sharing of malware, infrastructure, and operational practices among Chinese threat groups, as well as the possibility of access transfer between different actors.

Based on private telemetry, we identified a large collection of victim organizations compromised using ScatterBrain-obfuscated ShadowPad. Between July 2024 and March 2025, this malware was used in intrusions at over 70 organizations across various regions globally, spanning sectors such as manufacturing, government, finance, telecommunications, and research. We assess that the threat actor primarily gained initial foothold in the majority of these organizations by exploiting an n-day vulnerability in CheckPoint gateway devices, which aligns with previous research on ShadowPad intrusions involving the deployment of ransomware.

Among the victims, we identified the previously mentioned IT services and logistics organization that was at the time responsible for managing hardware logistics for SentinelOne employees. Victim organizations were promptly informed of intrusion specifics, which were swiftly investigated. At this point, it remains unclear whether the perpetrators' focus was solely on the compromised organization or if they intended to extend their reach to client organizations as well.

A detailed investigation into SentinelOne's infrastructure, software, and hardware assets found no evidence of secondary compromise. Nevertheless, this case underscores the fragility of the larger supplier ecosystem that organizations depend upon and the persistent

threat posed by suspected Chinese threat actors, who continuously seek to establish strategic footholds to potentially compromise downstream entities.

SentinelLABS will share a detailed public release on this topic in due course, providing further technical information on these activities, including observed TTPs, malware, and infrastructure.

Lessons Learned While Hardening Our Operational Ecosystem

Our analysis of the PurpleHaze cluster, and more specifically the potential indirect risk introduced via compromised third-party service providers, has reinforced several key insights around operational security and supply chain monitoring. Even when our own infrastructure remained untouched, the targeting of an external service provider previously associated with business logistics surfaced important considerations.

One immediate reminder is the necessity of maintaining real-time awareness not only over internal assets but also over adjacent service providers—particularly those with past or current access to sensitive employee devices or logistical information. When incidents occur near your supply chain, don't wait for confirmation of compromise. Proactively trigger internal reviews of asset inventories, procurement workflows, OS images and onboarding deployment scripts, and segmentation policies to quickly identify any exposure pathways and reduce downstream risk.

This leads to several defense recommendations:

- **Distribute Threat Intelligence Across Operational Stakeholders**
  Organizations should proactively share campaign-level threat intelligence with business units beyond the traditional security org—particularly those managing vendor relationships, logistics, and physical operations. Doing so enables faster detection of overlap with compromised third parties and supports early reassessment of exposure through external partners.
- **Integrate Threat Context Into Asset Attribution Workflows**
  Infrastructure and IT teams should collaborate with threat intelligence functions to embed threat-aware metadata into asset inventories. This enables more responsive scoping during incident response and enhances the ability to trace supply chain touchpoints that may be at risk.
- **Expand Supply Chain Threat Modeling**
  Organizations should refine their threat modeling processes to explicitly account for upstream supply chain threats, especially those posed by nation-state actors with a history of leveraging contractors, vendors, or logistics partners as indirect access vectors. Tailoring models to include adversary-specific tradecraft enables earlier identification of unconventional intrusion pathways.

While attribution continues to evolve and victim impact remains diverse, one thing is clear: well-resourced threat actors are increasingly leaning on indirect routes into enterprise environments. Investigations like this help us sharpen our defenses—not just around traditional digital perimeters but around the full operational footprint of our organization.

The Strategic Value of Cyber Threat Intelligence

In today's threat landscape, threat intelligence has evolved from a niche function into an essential pillar of enterprise defense—particularly for private sector organizations operating in the security space. As threat actors increasingly target security vendors for insider access, abuse of legitimate channels, and supply chain infiltration, the role of CTI in anticipating and disrupting these tactics has become more critical than ever.

One of the most tangible examples of this value is in internal talent acquisition and insider threat defense. Intelligence has become a frontline asset in identifying attempts by North Korean IT workers and other state-backed operatives to embed themselves in organizations under false pretenses. By flagging suspicious applicant patterns, cross-referencing alias histories, and tracking known tradecraft, CTI teams help hiring managers and HR avoid potential insider incidents before they start.

Our CTI capabilities must also directly support sales and channel operations. As criminal groups increasingly impersonate legitimate businesses to acquire security products through trusted resellers, intelligence plays a key role in verifying customer legitimacy and identifying anomalous purchase behaviors. By integrating intelligence insights into pre-sale vetting workflows, a crucial layer of protection is helping to ensure adversaries cannot simply "buy" their way into our technology stack.

Internally, threat intelligence informs and enhances how we defend our own technology and supply chain against highly targeted APT activity. From understanding how adversaries reverse-engineer our software to uncovering which parts of our technology stack they seek to compromise, CTI enables proactive hardening, smarter telemetry prioritization, and meaningful collaboration with product and engineering teams. In essence, intelligence acts as an early-warning system and a strategic guide—ensuring our defenses stay one step ahead of evolving threats.

Across every function—whether it's HR, Sales, Engineering, or Security—cyber threat intelligence is no longer a backroom function. It's embedded in the fabric of how we defend, operate, and grow as a business.