



Q2 / 2026

Threat reality report

Turn the latest intelligence into action

Reversesec

Stuart Morgan

Daniel Ashhami

Leonidas Tsaousis

Leo Stavliotis

Orpheus

Tim West

Sammar Smesme

A joint operation with



ORPHEUS

REVERSESEC

Contents

How this report works	3
Executive summary	4
Q1 2026 Threat intelligence	5
Cybercrime and ransomware	5
APT (Advanced Persistent Threat)	10
Hacktivism	14
Cross-cutting analysis	19
TTP analysis: VoidLink malware framework	19
Vulnerability landscape review	23
Sector threat levels	24
Q2 Outlook	25
Threat actors	25
AI as an external driver	28
Expected techniques	30
Testable real-world scenarios	34
Scenario A	34
Scenario B	36
Scenario C	38

From intelligence to testing

How this report works

Reversesec, an offensive-driven cybersecurity consultancy built on 30 years' experience, has partnered with **Orpheus**, a leading threat intelligence firm, to deliver an assessment of cyber threats that combines threat intelligence analysis with real-world red team insight. The report is a free resource for CISOs and senior leaders to support informed security decisions.

Translate intelligence into action

The report connects observed adversary behavior to testable models already used in red team and purple team exercises.

Observe current signals

Threat intelligence focuses on live activity and prominent techniques used by cybercrime, APT, and hacktivist groups.

Anticipate likely techniques

The Q2 outlook projects the techniques most likely to shape activity next quarter.

Prioritize testing

Sector-level overviews, TTP analysis and a vulnerability landscape review help security leaders decide what to test first.

Executive summary

This report links tactical intelligence with attack-chain simulation to show how threats are evolving, how they are likely to materialize, and where to focus defense. It covers Q4 2025–Q1 2026, with a forward view into Q2 2026. Threat intelligence is structured into three tracks: Cybercrime, APT, and Hacktivism.

The Bank of England's January 2026 CBEST thematic assessment found that simulated nation-state and organized criminal actors compromised financial firms by exploiting weaknesses in identity and access management, network segmentation, detection and response, and credential security.

Initial access was most often gained through social engineering, phishing, and third-party or supply chain compromise. While threat intelligence aided operational response, strategic integration and governance were less mature. Many attack paths reflected systemic weaknesses rather than novel exploits. Detection and response gaps delayed detection.

The report builds on those findings with analysis of the live threat landscape, showing how the same structural weaknesses are being used by financially motivated cybercriminals, state-sponsored actors, and hacktivist groups. Real-world threat-led emulation scenarios translate this intelligence into practical, testable adversary models to support proactive resilience validation.

How the report is organized: 1) identify systemic weaknesses exposed through controlled adversary simulation; 2) map observed activity that exploits those weaknesses at scale; 3) analyze TTPs and review the vulnerability landscape; 4) project the techniques most likely to dominate Q2 2026; 5) produce testable, real-world scenarios for resilience validation.

Recommendations: To strengthen resilience against advanced and persistent threats, organizations should test not only perimeter controls, but also identity governance, privileged access pathways, software supply chain integrity, and detection depth across hybrid and cloud-native environments




Part 1

Q1 2026

Threat intelligence

Cybercrime and ransomware

The charts below, based on Orpheus' internal reporting, illustrate which cybercrime threat each sector is most exposed to: data theft, ransomware, social engineering, or supply chain compromise. They are based on Orpheus' continual monitoring of the threat landscape.



In Q1 2026, cybercriminal threats remain high across sectors, with sustained data leak extortion and ransomware campaigns. Targeting remains primarily opportunistic and financially motivated, with an emphasis on targeting sectors that have a low tolerance for downtime, such as manufacturing and healthcare. January saw high activity from Qilin, Everest, and Akira ransomware franchises in the telecommunications, manufacturing, and healthcare sectors.


Everest is expanding its operations, increasingly acting as an Initial Access Broker (IAB). Cybercriminal and ransomware threat actors have maintained a focus on supply chain compromise and social engineering techniques for initial access. Orpheus has observed elevated service disruption to maximize impact and ransom negotiations through the targeting of overlooked physical access control systems and essential services.

Development 1: Mass exploitation of shared platforms enables extortion at scale

CLOP are currently one of the most prolific extortion actors, undertaking exploitation of widely deployed enterprise software and shared services platforms. While CLOP historically operated a ransomware-as-a-service (RaaS) model, since 2023, activity has increasingly focused on mass exploitation of internet-facing file transfer, enterprise resource planning, and collaboration platforms to enable downstream compromise of multiple organizations through a single vulnerability. Essentially, they have flipped the ransomware paradigm from 'quantity not quality' (targeting as many networks as possible) to 'quality not quantity'—undertaking a small number of campaigns that allow for efficient data theft from many third-party organizations.

This model was established through the exploitation of MOVEit Transfer and has since expanded to additional enterprise technologies, including late 2025 campaigns targeting Oracle E-Business Suite and Gladinet CentreStack. This reflects a continued investment in vulnerability discovery and rapid weaponization of newly disclosed or zero-day vulnerabilities.

Victims posted to CLOP's dark web data leak site (DLS) in recent months include large enterprise organizations across multiple sectors and geographies, with several victims experiencing data theft volumes exceeding one terabyte (claimed by CLOP). Recent victim postings have occurred in batches, often grouped by geography or campaign wave, consistent with compromise of shared enterprise platforms deployed across multiple organizations in the same region or supply chain. As targeting is largely platform or service-driven rather than sector-specific, it results in indiscriminate victimology across financial services, manufacturing, professional services, technology, and public sector organizations where vulnerable/compromised services are in use.



Targeting is driven by exploitable platforms and services, not by industry sector.

This is a highly efficient operational model that enables CLOP to maximize return on investment by exploiting a single platform to gain access to numerous downstream victim environments and applying extortion pressure at scale through DLS publication and direct victim notification.

Analysis:

CLOP's focus on enterprise software exploitation reflects attacker opportunity with the growing complexity of modern information system architecture, where organizations increasingly rely on interconnected off-premises infrastructure, cloud platforms, and software-as-a-service (SaaS) providers to support core operations. The widespread adoption of shared services and externally hosted enterprise platforms creates concentration risk, where compromise of a single provider or software vulnerability can enable access to multiple downstream organizations.

At the same time, increased availability of vulnerability research tooling, automated exploitation frameworks, and AI-assisted vulnerability discovery and exploitation techniques may accelerate identification, weaponization, and proliferation of exploitable conditions within the cybercriminal ecosystem. These conditions may increasingly enable financially motivated threat actors to conduct high-impact, multi-victim compromise through a small number of technically focused campaigns.

Development 2: Voice-based phishing captures SSO credentials

In late January 2026, ransomware collective Shiny Hunters (tracked as UNC6040) resurfaced with a new DLS listing major enterprises in the tech, education, retail, financial, and food service sectors. While the DLS is branded as "ShinyHunters", it is likely linked to activities of the broader extortion franchise tracked as Scattered Lapsus\$ Hunters (SLH).

ShinyHunters/SLH claimed responsibility for the recent Okta single sign-on (SSO) credential theft campaign, which was reportedly used to gain initial access to Crunchbase, Betterment, Match Group, and numerous other enterprise networks. Okta has warned¹ that adversaries are leveraging custom voice-based phishing (vishing) kits to steal SSO credentials for data theft from Okta, Microsoft, and Google accounts. Google Threat Intelligence (GTIG) reported² a wave of recent cybercriminal activity relying on sophisticated vishing tactics in which threat actors place phone calls to victim organizations' employees to trick them into visiting malicious credential harvesting sites and granting access.

¹ <https://www.okta.com/blog/threat-intelligence/phishing-kits-adapt-to-the-script-of-callers/>

² <https://cloud.google.com/blog/topics/threat-intelligence/voice-phishing-data-extortion>

Targeted employees typically have access to sensitive IT systems or internal tools and are lured by phone calls impersonating internal IT or helpdesk staff. Vishing is an increasingly popular tactic in social engineering campaigns abused by threat actors with varying capabilities, from malicious insiders to organized cybercriminal groups³.

Analysis:

Sophisticated social engineering, vishing, and credential theft align with the commonly used TTPs of Shiny Hunters/SLH. New phishing infrastructure affiliated with SLH has been recently identified that involves a “live phishing panel” that allows threat actors to perform man-in-the-middle attacks on login sessions in real-time to capture credentials and multifactor authentication tokens for SSO platforms, like Okta. This infrastructure is reportedly being leveraged to target over 100 high-value organizations since January 2026, demonstrating the continued effectiveness of social engineering campaigns against even mature organizations, and the continued criminal investment into technology underpinning such attacks. It is highly likely that this will be a persistent attack vector throughout 2026.


January 2026

Top 10 impacted sectors

1. Manufacturing
2. Technology
3. Healthcare
4. Professional services
5. Financial services
6. Consumer services
7. Construction
8. Education
9. Agriculture and food production
10. Public sector

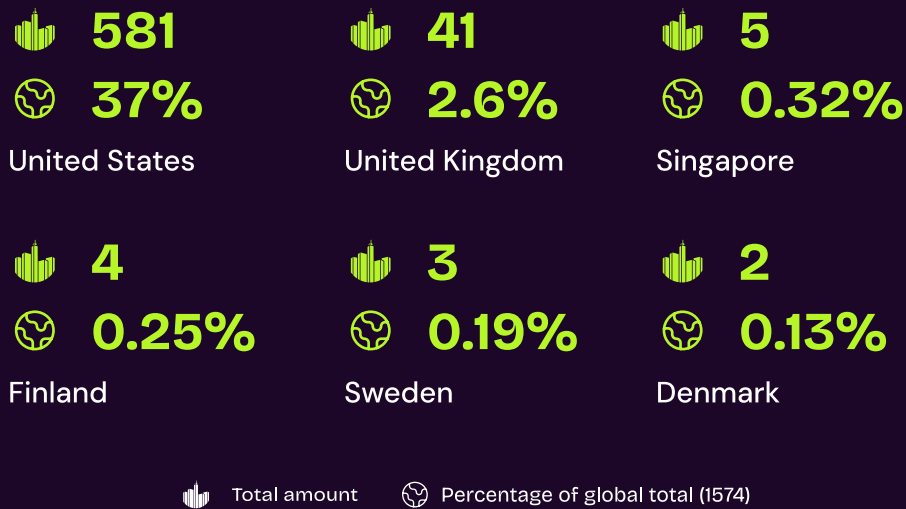
2025 - 2026

Monthly event count



³ <https://reliaquest.com/blog/threat-spotlight-shinyhunters-fast-tracks-saas-access-subdomain-impersonation/>

January – February 2026

Ransomware victims by country**Simulating cybercrime
in Reversesec's exercises**

As financially motivated threat groups pose priority threats to Reversesec's customers, their tactics and techniques have been simulated in exercises of different formats, to suit different needs.


The majority of "Attack Path Mapping" exercises conducted by Reversesec have included a ransomware deployment objective, following stakeholder requests. This has been interpreted at times as mass deployment and execution of encryptor software – optionally after exfiltration for "double extortion" scenarios, or in availability terms through widespread system disruption. These follow appropriate attack positioning, which in practical terms involves acquiring privileged access in cloud identity planes such as Entra ID, which is then abused to leverage MDM technologies such as Intune.

Reversesec has found that identifying and proving such paths consistently within limited timeframes, particularly in cloud infrastructure, requires relaxing of other constraints such as exercise stealth and environment knowledge. In some cases, there is an opportunity to covertly investigate the extent to which financial entities, who are required to undergo regulatory driven threat intelligence led penetration testing, are resistant to ransomware deployment. This is usually tested by demonstrating the ability to execute code across an estate which can modify data, but has included manipulation of encryption keys in the cloud.

APT (Advanced Persistent Threat)

Large institutions such as global banks, as well as critical infrastructure providers constitute the primary targets of APT groups. As a result, Reversesec can confirm that it's these threats that customers in such sectors approach to simulate in resilience building exercises, as previously identified by in-house Cyber Threat Intelligence teams.

The charts below illustrate, based on Orpheus' internal reporting, which nation-state each sector is most exposed to between China, Russia, Iran, and North Korea, based on their relative targeting intensity as observed by Orpheus.



In Q1 2026, most APT groups are refining existing strategies rather than developing new ones. North Korea added AI-generated malware development to its toolkit, further facilitating DPRK operations against financial and professional services sectors. Targeted campaigns against government entities and the public sector remain the greatest nation-state threat globally, with increased targeting of US government institutions by China-linked threat groups. Social engineering, especially targeted spear-phishing to deliver existing and novel malware, and vulnerability exploitation, remain common TTPs for state-sponsored actors.

However, other means of compromise are becoming increasingly common. Most organizations are investing in training on older techniques, and the success rate of anonymized external phishing is declining. SMS phishing (smishing) is gaining attention, mainly for credential harvesting attacks. Threat actors typically register phone numbers and send text messages using pretexts similar to phishing emails. They request that users authenticate to a service to confirm access, or they demand it under threat of financial penalties or loss of access.

In general, high volume generic phishing campaigns which require a user to open or execute an untrusted attachment are unsuccessful compared to highly targeted multi-modal campaigns. For example, establishing a dialogue using Teams or Zoom chat to establish legitimacy, followed by requesting that a document is opened, are techniques which are more successful.

Development 1: React2Shell rapidly weaponized across multiple sectors and regions

Toward the end of Q4 2025, espionage groups linked to China, Iran, and North Korea, and financially motivated threat actors, exploited a React Server Components (RSC) flaw shortly after it was disclosed on 3 December 2025.

The maximum-severity vulnerability, tracked as CVE-2025-5182, dubbed "React2Shell", is a critical unauthenticated remote code execution (RCE) flaw that stems from how React decodes payloads sent to Server Function endpoints. Improperly validated or maliciously crafted HTTP requests triggered unsafe deserialization, leading to unauthenticated remote code execution, even where the application does not explicitly use custom server functions. React2Shell affected a significant number of systems due to the prevalence of RSC, a feature of the React JavaScript library, in widely used frameworks, including Next.js, which is adopted in many well-known services such as Netflix and TikTok.

Various threat clusters quickly abused the RSC security gap to conduct diverse campaigns against multiple sectors and regions. Within hours of React2Shell's disclosure, at least five China-linked espionage groups deployed the following payloads via the exploit: UNC6600 delivered a MINOCAT tunneler; UNC6586 executed a SNOWLIGHT downloader; UNC6603 deployed an updated HISONIC backdoor⁴. Further attribution was attained through shared and reused infrastructure associated with China-affiliated UNC5454 (Earth Lamia) and Jackpot Panda, as well as undisclosed Iran-nexus actors. Days after the China-sponsored exploitation, North Korean actors leveraged React2Shell to deploy a previously undocumented remote access trojan⁵ EtherRAT, exhibiting tooling

⁴ <https://cloud.google.com/blog/topics/threat-intelligence/threat-actors-exploit-react2shell-cve-2025-5182>


⁵ <https://thehackernews.com/2025/12/north-korea-linked-actors-exploit.html>

overlap with the DPRK's Contagious Interview campaign targeting blockchain and Web3 developers.

Observed post-exploitation behaviors included deployment of the XMRig cryptocurrency miner obtained from GitHub, reconnaissance data collection, environment variable scanning, retrieval of secrets, raw SSH shell creation, botnet malware deployment, and general reconnaissance and vulnerability scanning, which likely includes researcher activity. After the initial disclosure of CVE-2025-55182, three additional React vulnerabilities were identified: CVE-2025-55183, CVE-2025-55184, and CVE-2025-67779 (CVE-2025-66478 was marked as a duplicate of CVE-2025-55182)⁶. On 15 December, the React2Shell exploits enabled a sophisticated credential-stealing campaign named "Operation PCPcat", which compromised over 59,000 Next.js servers worldwide to harvest sensitive authentication data at an industrial scale⁷.

Assessment:

At the time of React2Shell's disclosure, Orpheus assessed its likelihood of future exploitation at 95 per cent and determined that it was almost certain that threat actors could outpace "time-to-patch" due to React's open-source technology. While patches were released immediately after disclosure, the raw code changes made in the patch were available for anyone to review and reverse engineer, increasing the chances that proof-of-concept (PoC) exploit code could become available to threat actors. Indeed, within hours, verified working PoC code was made public and quickly leveraged by state-aligned groups as well as cybercriminals seeking to monetize access via cryptocurrency mining.



APTs rapidly weaponize newly disclosed vulnerabilities, often within hours, integrating public exploits to drive large scale, multi-CVE exploitation.

The swift, mass exploitation of React2Shell demonstrates APTs' high capability and intent to monitor for new vulnerability disclosures and rapidly integrate public exploits into their scanning infrastructure to compromise systems at scale, targeting multiple CVEs simultaneously to maximize exploitation success. In one case, a threat cluster was observed attempting multiple exploit payloads, totaling 116 requests in 52 minutes, while systematically

⁶ <https://cloud.google.com/blog/topics/threat-intelligence/threat-actors-exploit-react2shell-cve-2025-55182>

⁷ <https://cybersecuritynews.com/new-pcpcat-exploiting-react2shell-vulnerability/>

troubleshooting⁸. This indicates persistent, methodical intrusion activity rather than opportunistic automated scanning, with actors actively refining techniques against live environments.

Development 2: Wiper tools disrupt Polish distributed energy sites

During Q1 2026, Polish authorities confirmed that, on 29 December 2025, coordinated cyber campaigns targeted over 30 wind and photovoltaic farms, a private manufacturing company, and a large combined heat and power (CHP) plant supplying heat to almost half a million customers in Poland.

The event, attributed to the Dragonfly APT group, was assessed by the Polish national Computer Security Incident Response Team (CERT Polska) as having a purely destructive objective, rather than simple reconnaissance or short-term disruption, especially as the attacks struck while Poland was facing dangerously low temperatures and snowstorms. Nevertheless, the impact was minimized to disruption rather than destruction, as the attacks failed to interfere with the heat supply. Around 30 distributed energy sites were affected, and in several cases, operators lost visibility into their systems and the ability to remotely manage equipment. Despite these impacts, electricity generation itself continued, and grid stability was maintained throughout the incident⁹.

Adversaries gained access through exposed or poorly secured systems, after which they moved into OT networks and interfered with remote terminal units and associated communications infrastructure. The activity appeared methodical, with similar techniques repeated across multiple sites, indicating prior knowledge of the environment and its underlying technologies.

Threat actors deployed two previously undocumented wiper tools: a native Windows binary identified as 'DynoWiper' and a PowerShell-based script dubbed 'LazyWiper', both designed to cause permanent damage and irreversible destruction of data, further supporting the suspected objective of deliberate destructive actions. Four DynoWiper variants were used against the renewable energy farms and the CPH plant to delete data, disrupt system functionality, and render IT and operational technology devices unusable, including industrial controllers and supporting infrastructure. Attackers targeted the manufacturing company with LazyWiper, which overwrites files on the system, rendering it unusable and irrecoverable. The overwriting process was implemented via a C# function named WriteRandomBytes that was likely generated using a large language model (LLM)¹⁰.

⁸ <https://aws.amazon.com/blogs/security/china-nexus-cyber-threat-groups-rapidly-exploit-react2shell-vulnerability-cve-2025-55182/>

⁹ <https://cert.pl/en/posts/2026/01/incident-report-energy-sector-2025/>

¹⁰ https://cert.pl/uploads/docs/CERT_Polska_Energy_Sector_Incident_Report_2025.pdf

While early reporting suggested the involvement of the Russian GRU-sponsored Sandworm threat actor, CERT Polska assessed that the tactics, targeting rationale, and operation focus overlap more with the long-established Russian FSB-aligned cyber espionage group, Dragonfly (also tracked as Static Tundra, Berserk Bear, and Ghost Blizzard). The Dragonfly activity cluster is associated with a strong interest in the energy sector, high capability to target industrial devices, and similar infrastructure leveraged in the campaigns against Poland, including compromised VPS servers and Cisco routers, routers, and traffic patterns that were used to obtain initial access, exfiltrate data, establish VPN tunnels for wiper malware deployment, and damage the server's RAID array disks¹¹.

Given the high technical competence of these actors, it is crucial to allow for assessments to simulate similar levels of competence by not artificially restraining red teams to only "known" TTPs, but rather by focusing on the skill level of the threat actors. The use of Generative AI technologies to streamline capability development also has a place in this, in line with documented¹² adoption of LLM-powered workflows by threat groups in various steps of the attack lifecycle.

Assessment:

This incident is notable because it focused on distributed energy resources rather than central grid control systems. As renewable energy continues to expand, these assets represent an increasingly attractive target for nation-state-level threat actors. Compromising multiple smaller sites can still produce strategic effects, particularly if such access is combined with timing or broader geopolitical objectives, as with this attack that took place during the period just before New Year's, when the Polish population was vulnerable to significantly low temperatures.

Energy sector organizations should treat this incident as a warning that distributed energy assets are now firmly within the scope of advanced state-sponsored cyber operations. While recent activity highlights the evolving capabilities of Russian FSB-linked groups like Dragonfly, other state-aligned actors, such as Sandworm, continue to target Western energy and industrial systems, demonstrating a persistent and broad threat landscape.

Hactivism

Hactivist activity relevant to the UK and EU is dominated by ideologically motivated distributed denial of service campaigns conducted by pro-Russian collectives targeting government,


¹¹ https://cert.pl/uploads/docs/CERT_Polska_Energy_Sector_Incident_Report_2025.pdf

¹² <https://services.google.com/fh/files/misc/advances-in-threat-actor-usage-of-ai-tools-en.pdf>

critical national infrastructure, and private sector organizations across the UK, EU, and US. As is typical, activity is closely aligned with geopolitical developments, particularly Western support for Ukraine, and was characterized by repeated volumetric attacks against public-facing web infrastructure such as websites, customer portals, and authentication services. Technical methods remained consistent, with attacks primarily leveraging UDP and TCP floods, SYN floods, HTTP-layer floods such as HTTP_Loris, and DNS amplification, often enabled through commodity botnets and publicly accessible projects such as DDoSIA.


Hactivist activity linked to the Israel-Iran conflict remained persistent during this period, with pro-Iranian and pro-Israeli collectives conducting website defacements, Distributed Denial of Service (DDoS) attacks, and data leak operations against government entities, critical infrastructure, and private sector organizations to signal ideological alignment, amplify political messaging, and retaliate against perceived adversaries.

The charts below, based on Orpheus' internal reporting, illustrate which cybercrime threat each sector is most exposed to: data leak, DDoS, defacement, or disruption.



Hactivist TTPs have remained relatively consistent over the last 12 months. Most DDoS attacks are easy to mitigate with the

correct technology in place, and pro-Russian hackers will target organizations without protection. This said, Cloudflare reported an increase in frequency and size of DDoS attacks in Q3 2025¹³. High-bandwidth or amplification attacks are generally not part of threat simulations or other professional assessments, due to the low sophistication and minimal value offered. There are specific sectors which place emphasis on denial of service, for example implantable medical devices or very heavily regulated stock trading, but other forms of attack are usually more prevalent than amplification attacks.



Targeting is broad and opportunistic, with threat actors favoring organizations with immature DDoS protection.

#OpUK launched in Q1 2026, a coordinated pro-Russian hacker campaign responding to UK support for Ukraine and NATO partners. It relies primarily on repeated low-complexity DDoS attacks against government, public services, and commercial organizations to generate disruption and visibility rather than lasting impact. Targeting is broad and opportunistic, with threat actors favoring organizations lacking mature DDoS protection, reinforcing OpUK's role as a messaging and pressure operation rather than a technically advanced cyber threat.

Development 1: Broad DDoS campaigns target European organizations

While recent government warnings relating to ICS/OT targeting are valid, in January 2026, the observable activity from named hacker campaigns (such as #OpUK) was largely focused again on web-layer disruption (DDoS) and opportunistic website outages of government, finance, and public-facing services, with no widely reported, confirmed ICS/OT breaches attributed to these groups in January 2026 specifically.

#OpBelgium and Denmark-focused campaigns

Pro-Russian hacker collectives conducted coordinated DDoS campaigns targeting Belgian and Danish government organizations, telecommunications providers, and public infrastructure in response to political statements and national policy positions. These operations also affected financial institutions and public sector services, with disruption primarily limited to public-

¹³ <https://blog.cloudflare.com/ddos-threat-report-2025-q3/>

facing websites and portals, while maintaining high visibility and propaganda amplification.

#OpFrance and broader European financial targeting

Hacktivist campaigns targeted French national digital infrastructure, including La Poste¹⁴, disrupting online banking, postal services, and customer-facing systems through sustained volumetric DDoS activity. These operations formed part of a wider European targeting of financial and public sector organizations using botnet-enabled traffic flooding to overwhelm web infrastructure.

#OpUK and associated European expansion campaigns, including #OpDenmark

Pro-Russian hacktivist collectives launched campaigns targeting UK local government, transport hubs, telecommunications, financial services, and commercial organizations, primarily using repeated DDoS attacks against websites, login portals, and public-facing services. Parallel activity targeted Danish, German, and Italian government and infrastructure organizations, while financial institutions, manufacturing bodies, and professional services firms were opportunistically targeted due to their association with national economic and critical infrastructure ecosystems.

These campaigns demonstrated the continued and regular use of pseudo-opportunistic targeting, with attackers selecting organizations based on geographic relevance and exposure of internet-facing services where impact is deemed to be possible. While specific geographic operations are common, targeting of Ukrainian organizations is also consistent.

Development 2: Hacktivists target prominent events

Hacktivist activity is also highly responsive to individual events prominent on a global scale. Pro-Russian hacktivist group NoName057(16) has repeatedly conducted distributed denial of service attacks against Swiss government agencies, infrastructure providers, and related organizations during the World Economic Forum. Furthermore, the Winter Olympics have consistently been targeted by hacktivist and politically motivated cyber actors due to their global visibility and geopolitical symbolism¹⁵ – continuing a precedent started in 2018 with “Olympic Destroyer” malware attacks. It is likely this hacktivist activity was coordinated with state-sponsored actors, as it appeared choreographed with physical disruption of transport hubs and online misinformation campaigns.

¹⁴ <https://www.lapostegroupe.com/fr/actualite/incident-informatique-du-22-decembre>

¹⁵ <https://www.bbc.co.uk/sport/articles/cqj25wyjx1no>

Across all sectors, hacktivist activity remained heavily concentrated at the network and application layer, with disruption largely limited to public-facing services rather than internal systems or operational environments. Local government and public sector organizations were among the most frequently targeted due to their visibility and public accessibility, while telecommunications providers experienced repeated attempts to disrupt customer-facing portals. Educational organizations experienced isolated defacement and intrusion incidents linked to political and social causes.

Overall, hacktivist campaigns during this period demonstrated sustained targeting frequency across multiple sectors, consistent use of volumetric denial of service techniques, and continued reliance on available infrastructure and botnet-enabled traffic generation to conduct disruptive operations. An Orpheus customer was included on the target list of a Pro-Russian hacktivist, and therefore, Orpheus was able to confirm the traffic was sinkholed, and the attack caused no impact. The continual DDoS activity against the UK and EU likely caused very low, or more likely, no impact.

Part 2

Cross-cutting analysis

TTP analysis: VoidLink malware framework

Security researchers recently identified VoidLink, an advanced malware framework with cloud and container environments capabilities, designed for long-term access to Linux systems. The framework is associated with Chinese-speaking developers and was likely designed to be productized and sold commercially to red-team operators¹⁶.


VoidLink provides operators with custom loaders, implants, rootkits, and plugins. Its architecture is flexible and modular, and centered around a custom Plugin API, enabling rapid adaptability to specific environments and objectives. It is a cloud-first implant designed to operate in modern infrastructure and recognize major cloud environments, currently AWS, GCP, Azure, Alibaba and Tencent. It is also able to detect when running inside Kubernetes or Docker and tailor its behavior accordingly. This capability is, in part, what made Cobalt Strike so popular with red teams and cybercriminals.

A threat actor, UAT-9921, active since 2019, has been targeting organizations in the technology and financial services sectors recently and has been observed utilizing VoidLink. The group comprises hosts and deploys VoidLink command-and-control infrastructure to establish persistent remote access. These compromised systems are then used to conduct network scanning activities, both laterally within internal environments and externally against internet-facing assets.

¹⁶ <https://research.checkpoint.com/2026/voidlink-the-cloud-native-malware-framework/>

TECHNIQUE	MITRE TTP ID	DESCRIPTION
Masquerading	T1036	Threat actors may disguise VoidLink binaries, plugins, or scripts as legitimate system or DevOps processes to evade detection and blend into normal Linux or cloud workload activity.
Deobfuscate or Decode Files or Information	T1140	Threat actors may decrypt or unpack VoidLink plugins and payloads at runtime to conceal functionality and reduce the effectiveness of static detection.
Credential Access: Cloud Account Credentials	T1556.001	Threat actors may harvest cloud credentials, API tokens, and secrets from compromised Linux systems to maintain access to cloud resources and expand control.
Lateral Movement: Remote Services	T1021	Threat actors may use harvested credentials to authenticate to remote systems or cloud services, enabling lateral movement without deploying additional malware.
Ingress Tool Transfer	T1105	Threat actors may download additional tools or plugins through VoidLink to extend postcompromise capabilities without redeploying the core implant.
Exfiltration Over Web Service	T1567.002	Threat actors may exfiltrate data over HTTPS or other web services to blend with legitimate outbound traffic and evade networkbased detection.
Indicator Removal on Host	T1070	Threat actors may delete logs, files, or other forensic artefacts to reduce visibility and hinder incident response.
System Network Connections Discovery	T1049	Threat actors may enumerate active network connections and listening services to understand the environment and identify opportunities for further compromise.
Modify Cloud Compute Infrastructure	T1578.002	Threat actors may alter cloud instance or container configurations to weaken security controls or maintain longterm access.

TECHNIQUE	MITRE TTP ID	DESCRIPTION
Data from Information Repositories	T1213	Threat actors may access source code repositories or configuration stores to collect sensitive data or enable supplychain compromise.



CATEGORY	BEHAVIOUR	DESCRIPTION
Advanced Linux Implant	VoidLink Implant Core	VoidLink operates as a sophisticated Linux commandandcontrol agent supporting over 100 command types, enabling longterm remote control, tasking, and postexploitation operations across cloud and container environments.
Rootkit Capabilities	KernelLevel Stealth Mechanisms	VoidLink contains embedded kernel module code such as vl_stealth.ko and ss_loader, enabling syscall hooking (getdents64), netfilter manipulation, and kretprobes to hide processes, network ports, and files from userspace monitoring.
Credential Harvesting	MultiSource Credential Access	VoidLink actively targets sensitive credentials including SSH private keys, AWS credentials, Kubernetes service account tokens, Git credentials, browser cookies, and environment variables to expand access and enable lateral movement.
Lateral Movement Enablement	Environment and Access Enumeration	VoidLink parses SSH known_hosts, detects Kubernetes and cloud environments, and performs automated capability assessment to identify opportunities for lateral movement or privilege expansion.
Covert Communication	Stealthy C2 Channels	VoidLink supports ICMPbased covert communication using a magic value (0xCODE), multimode encryption, and a custom binary C2 protocol to evade networkbased detection.
AntiForensic Activity	Artefact and Log Manipulation	VoidLink includes functionality to clear logs, wipe command history, timestamp files, and selfdelete implants and plugins to hinder forensic analysis and incident response.

CATEGORY	BEHAVIOUR	DESCRIPTION
Modular Architecture	Dynamic Plugin System	VoidLink uses a dynamic ELF loading mechanism that allows operators to deploy additional plugins at runtime, extending functionality without modifying the core implant.
Evasion Techniques	AntiAnalysis and Camouflage	VoidLink is statically linked with stripped symbols, employs antireverse engineering modules, and uses a camouflage executor to disguise execution context and resist static and dynamic analysis.

In line with widely accepted notions of the detection engineering practice¹⁷, Reversesec believes that focusing efforts on specific campaigns and tools will provide limited value, at risk of becoming dated before the exercise even concludes.

Instead, through the atomic “Attack Detection Capability Assessment” approach, purple team exercises can simulate multiple campaigns from multiple threat groups at once, by isolating TTPs in individual attacker actions. With each such “test case” then examined independently in terms of Log collection, Alerting, Prevention metrics, dense data can be collected and analyzed, allowing extraction of insights such as:

- The <example> solution recorded the majority of detections raised overall
- Visibility was higher in Windows systems than in Unix servers
- Exfiltration test cases were largely unprevented

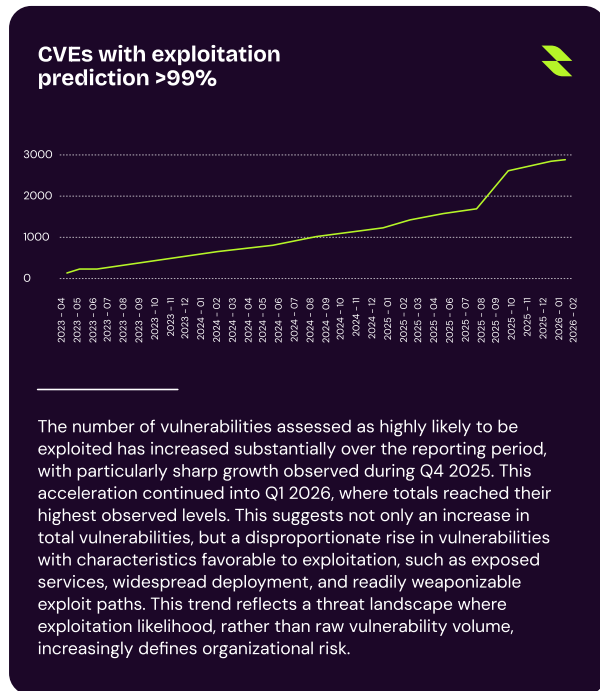
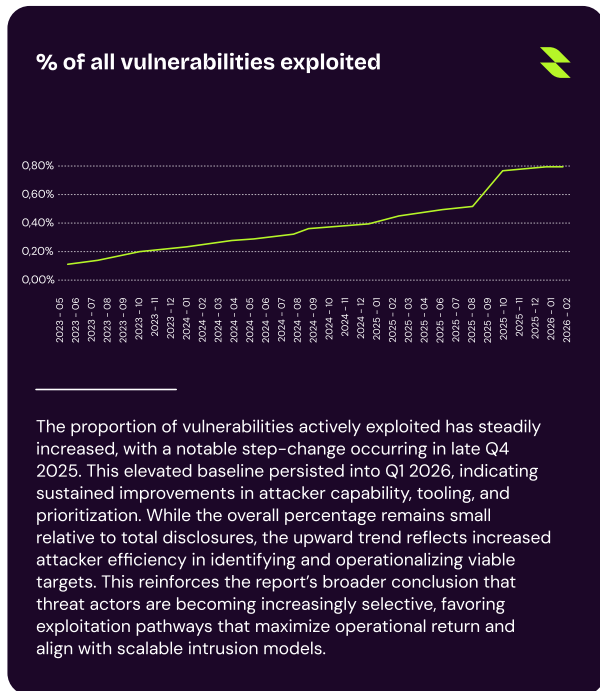

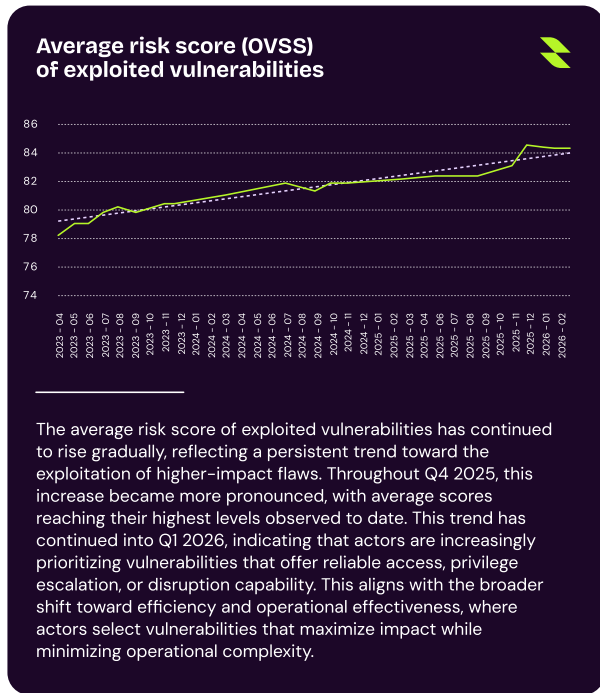
Malware framework emulation

Nevertheless, in notable instances, Reversesec has been engaged to reproduce, and safely detonate malware chains with as close likeness to original frameworks as possible, while maintaining operational guardrails for a professional setup. This has included scripts, compiled binaries, as well as identical artifacts such as ransomware notes, which were engineered based on publicly available documentation, aiming to trigger the same indicator-based alerts such as YARA rules and antivirus signatures. Through this approach, teams have been able to model detection and response capabilities – with a focus on technologies – against a specific threat actor campaign of interest.

¹⁷ <https://detect-respond.blogspot.com/2013/03/the-pyramid-of-pain.html>

Vulnerability landscape review

Software vulnerabilities remain a primary initial access vector for a plethora of threat actors, from APT threats to well-resourced cyber criminals. For many organizations, strong cyber hygiene is the best mitigation against opportunistic threats. Orpheus monitor the vulnerability risk environment and track its development using proprietary metrics proven to be more accurate than commonly used open-source metrics (CVSS, EPSS). The following graph depicts this evolving landscape, which is vital when attempting to forecast the environment into the next quarter.



Sector threat levels

The following is based on Orpheus' analysis of the current threat level, per sector. Of the sectors included, finance is the most targeted, facing high-level threats from APTs and financially motivated cybercriminals, and a moderate threat level from hacktivists. Telecommunications, the public sector, healthcare and manufacturing face high-moderate threats from cybercriminals, while manufacturing is at slightly lower risk of APT attacks. Cybercrime poses the highest and hacktivism the lowest overall. The following charts show a high-level threat score from the three main threat actor types (APT, Cybercriminal, Hactivist). Also included for each sector is a list of observed threat actor objectives.



Part 3

Q2 Outlook

Threat actors

Cybercrime and ransomware:

At the time of writing this report, the seven-day rolling average has increased by 38 per cent over the preceding week period¹⁸. Over the preceding six months, the average number of monthly victims posted to leak sites has also continually increased. There is typically a seasonality to ransomware victim numbers, with the second quarter of each year often busier than the first.

Looking into Q2 2026, ransomware activity is therefore likely to remain elevated, with opportunistic targeting aligned to both seasonal trends and broader geopolitical instability. Criminal groups will continue to exploit the same weaknesses observed throughout 2025, including exposed edge devices, unpatched critical vulnerabilities, and identity-layer misconfigurations. Objectives are unlikely to change significantly.

Financial extortion remains the primary driver, with actors seeking to maximize payment probability through data theft, operational disruption and reputational pressure via leak sites. The continued integration of ransomware operations into broader cybercriminal service ecosystems, including IABs and RaaS affiliates, will maintain high attack volume and rapid victim turnover. Unless disrupted by major law enforcement action or internal fragmentation within dominant groups, Q2 2026 is expected to sustain or exceed current posting volumes.

Tactically, ransomware tradecraft is now extremely broad and mature. Operators routinely blend commodity malware, living-off-the-land techniques, credential harvesting, vulnerability exploitation and layered social engineering to gain and expand access. The success of Scattered Spider in leveraging vishing campaigns against IT help desks demonstrates that human-enabled intrusion remains highly effective, particularly where identity reset processes can be manipulated¹⁹. This approach is

¹⁸ Orpheus Ransomware victim stats

¹⁹ <https://www.cisa.gov/news-events/cybersecurity-advisories/aa23-320a>

likely to continue into Q2 2026 and beyond. Similarly, consent phishing and OAuth abuse remain under-recognized by users, who often do not understand the implications of granting token permissions to malicious applications.

These techniques are not emerging trends specific to Q2 2026, but rather persistent features of the ransomware ecosystem throughout 2025 and into 2026. As identity becomes the primary control plane in modern enterprise environments, many ransomware operators will continue prioritizing social engineering and authentication bypass over purely technical exploitation, widening the effective attack surface regardless of perimeter security maturity.

APT:

Looking ahead into Q2 2026, APT campaign objectives are unlikely to change materially unless driven by a significant geopolitical shock. Strategic priorities will continue to center on intelligence collection against government, defense and diplomatic entities, persistent access into critical infrastructure, and revenue generation to offset sanctions pressure. China-linked groups are highly likely to sustain long-term espionage operations against US and allied public sector institutions (particularly telecommunications organizations), while Russian actors will continue pre-positioning within European energy, logistics and industrial networks, particularly where distributed assets create asymmetric leverage. The rapid exploitation of React2Shell illustrates that state-sponsored actors are now structurally integrated into vulnerability disclosure cycles, operationalizing public proof-of-concept code within hours.

This pattern is expected to persist into Q2, with APT groups prioritizing scalable exploitation, multi-CVE chaining and access persistence rather than immediate disruption. However, in the event of a major geopolitical escalation, these access operations could quickly transition from intelligence collection to destructive or coercive activity. North Korean operators are increasingly professionalizing advanced social engineering, using AI-enhanced persona development, deepfake-enabled recruitment lures and tailored spear-phishing to gain access to software development environments. Their operations are evolving beyond standalone financially motivated campaigns and are becoming embedded within elements of the RaaS ecosystem, broadening their targeting aperture considerably.

By leveraging RaaS infrastructure and affiliate-style partnerships, DPRK actors can blend state objectives with criminal distribution models, increasing scale while obscuring attribution. In Q2 2026, it is highly likely they will continue prioritising software supply chain compromise through developer impersonation, malicious

dependency injection and CI/CD pipeline infiltration, particularly in Web3, fintech and SaaS environments.

Performing assessments from the starting point of a compromised developer is crucial for uncovering and mitigating these kinds of attack paths. These assessments should be broad in scope; if developers have access to CI/CD pipelines, then so do attackers. It is very common for organizations to perform risk assessments on objectives and, even if a compromised developer is a viable threat, avoid proceeding with those scenarios. There is always a margin for error in such operations, and due to stability risk, organizations are hesitant to proceed.

Nevertheless, the benefits of such scenarios outweigh the drawbacks, and they should be allowed to run. There are many workaround solutions that would still be beneficial to companies. A commonly adopted approach is for the red team to position itself in the network with sufficient permissions to compromise a pipeline or push code to a development platform such as GitHub. In such situations, the organization still covers a compromised developer scenario and damage control is achieved.

Furthermore, some in-house development platforms run with excessive privileges and are misconfigured which allows threat actors to target them. If such a compromise is successful, then high privileged access has been obtained which undoubtedly benefits the attacker. For example, if it is possible to access a Jenkins console, it is usually possible to execute code on the underlying server or manipulate the CI/CD pipeline.

Hacktivism:

Hacktivist activity relevant to the UK and EU will remain a persistent yet largely low-impact element of the threat landscape in Q2 2026. Reporting around the February 2026 Winter Olympics highlights that hacktivist communities escalated their chatter and coordination ahead of the Games and leveraged the event's global visibility to publicize grievances.

Looking ahead to Q2 2026, hacktivist operations are expected to remain frequent and opportunistic. Pro-Russian collectives will likely continue to direct DDoS and website disruption campaigns against Western governments and financial institutions in response to ongoing support for Ukraine and other geopolitical tensions. The 2026 FIFA World Cup, which begins on 9 June, will be an attractive target; the tournament's scale and visibility make it a "marquee target" for anti-Western hacktivists, raising the risk that ticketing, broadcasting and sponsorship systems could be disrupted. While defacement and DDoS attacks are likely to persist, there is a possibility that hacktivist campaigns may attempt higher-impact actions such as data leaks or supplychain exploitation.


Anti-Israel and pro-Iran hacktivist activity remains persistent; however, if the United States increases its involvement in Iran or conducts kinetic operations in the region in Q2 2026, this ecosystem will likely escalate disruptive hacktivist campaigns against Western organizations, including DDoS attacks, website defacements, data leaks, and “fake ransomware” or wiper operations designed primarily for psychological and reputational impact rather than financial gain.

Nevertheless, the predominant threat in Q2 will remain at the network and application layers, and organizations with robust DDoS mitigation and incident response plans should be able to manage the risk.

AI as an external driver

Artificial intelligence is now almost certainly viewed as a strategically critical capability, underpinning economic productivity and national power. Geopolitical competition increasingly centers on access to rare earth minerals and advanced semiconductor supply chains essential to continued AI development.

However, competition is no longer confined to hardware. AI models themselves are strategic assets. Market reactions illustrate their economic weight. The release of DeepSeek in early 2025 disrupted global technology stocks. In Q1 2026, the launch of new Claude security-focused models reportedly erased billions from cybersecurity company valuations. Advanced AI capability is therefore directly influencing market stability and sector confidence²⁰. In Q1 2026, Anthropic disclosed details of what it described as an advanced distillation attack targeting its models. While the company stopped short of formal attribution, it explicitly referenced Chinese-developed models, including DeepSeek, Moonshot and MiniMax, as being used in attempts to illicitly extract Claude’s capabilities to enhance their own systems²¹.



Removing safeguards increases the risk that powerful AI capabilities are repurposed for offensive cyber operations.


²⁰ <https://www.cnbc.com/2026/02/23/cybersecurity-stocks-anthropic-ai-crowdstrike.html>

²¹ <https://www.anthropic.com/news/detecting-and-preventing-distillation-attacks>

The framing strongly implied concern regarding Chinese state-linked or state-tolerated activity, particularly given the strategic value of frontier AI systems. Anthropic further warned that distilled models may lack enforced safety guardrails. The removal or absence of such safeguards increases the risk that powerful AI capabilities could be repurposed for offensive cyber operations, disinformation campaigns, surveillance applications, or other national security use cases.

Beyond state actors, the proliferation of unguarded high-capability models presents clear criminal implications. It has long been assessed that sufficiently capable AI systems could accelerate ransomware operations, fraud, and large-scale social engineering. As models become more capable and more widely accessible, barriers to entry for sophisticated cyber activity are likely to lower. The AI landscape continues to evolve at pace and precise forecasting remains challenging. However, it is almost certain that state, proxy and criminal threat actors will increasingly integrate advanced AI capabilities into their operations over the coming quarters, with strategic competition between the United States and China remaining a central driver.

Attacks against AI models in production environments are emerging as a distinct and growing threat category. As organizations embed LLMs and related systems into core business workflows, customer services, and security tooling, these models become high-value operational assets. This materially expands the attack surface.



Threat actors increasingly target AI systems embedded in critical business functions.

Many AI attack vectors remain at the proof-of-concept stage, largely demonstrated in academic settings such as model manipulation or training data poisoning. However, reporting from Orpheus and Reversesec indicates that threat actors are increasingly targeting AI systems as they become embedded within critical business functions. Activity observed across Q4 2025 and Q1 2026 suggests a shift from theoretical risk to operational targeting.

Notable recent developments include:

- The emergence of reprompt attacks, where malicious prompts were concealed within URL parameters, such as the “q” field, to influence model outputs²².

²² <https://www.varonis.com/blog/reprompt>

- Malicious Chrome extensions masquerading as AI assistant tools in a campaign referred to as Prompt Poaching²³.
- A software supply chain attack targeting the publishing process of the widely used AI assistant Cline. Attackers compromised publishing tokens to distribute a malicious package, which executed a post-install script, silently deploying the OpenClaw agent onto affected systems. This created unauthorized, autonomous access within developer environments²⁴.

Throughout Q2 2026, it is highly likely that attacks against common AI services and tooling will increase in frequency and sophistication. Both security researchers and threat actors are actively exploring novel attack vectors, accelerating the evolution of AI-focused tradecraft.

Expected techniques

Continuing from the content put forward in this product, particularly the threat actor forecast of Q2 2026, the following techniques are assessed to be heavily leveraged in Q2 2026 and should be a focus for organizations seeking to proactively test their resilience to initial access vectors expected to be effective, damaging and common across the full threat actor spectrum in Q2 2026.

Phishing

- Expect more SSO authorization phishing (i.e. device code/consent) that tricks employees into granting threat actors cloud access without entering credentials, yielding sessions usable across the enterprise's cloud services, including mail, files, and collaboration. Threat actors will continue leveraging and refining techniques to obtain enterprise access that appears legitimate, accelerating SaaS lateral movement and data theft.
- Phishing will likely focus more on accounts payable/accounts receivable (AP/AR) and payroll workflows by hijacking trusted email threads, including via supplier mailbox compromise. Wire transfer business email compromise (BEC) attacks increased 136 per cent from Q3 to Q4 2025²⁵, indicating that financial workflows, especially AP/AR and vendor payment changes, are an increasingly valuable phishing target for financially motivated threat actors compared to 2025. A seasonal phishing spike can be expected in Q2 2026 as phishing attacks rose 12 per cent (mainly targeting the financial/payment, SaaS/Webmail, and retail/e-commerce sectors) and wire transfer BEC attacks rose 27 per cent from Q1 to Q2 in 2025²⁶.

²³ <https://www.blackfog.com/prompt-poaching-fake-chatgpt-extensions/>

²⁴ <https://www.endorlabs.com/learn/supply-chain-attack-targeting-cline-installs-openclaw>

²⁵ https://docs.apwg.org/reports/apwg_trends_report_q4_2025.pdf

²⁶ https://docs.apwg.org/reports/apwg_trends_report_q2_2025.pdf

- As the number of known phishing kits doubled in 2025, sustained growth of phishing industrialization and dominance of phishing-as-a-service (PaaS) is expected to continue in 2026²⁷. Expect more adversary-in-the-middle, kit-driven phishing that steals session tokens/cookies and then pivots from one service (e.g., Microsoft 365) into broader SaaS and identity brokerage. High-volume campaigns leveraging PaaS kits will continue to rely on URL obfuscation, CAPTCHA abuse, malicious QR codes, and MFA bypass, increasing the risk of identity compromise.
- Expect more phishing aimed at developers and platform engineers to steal repository and registry access and API keys, enabling malicious packages (and, increasingly, trained LLM models) which can cascade into software releases and AI services enterprise-wide.

Vishing

- Vishing, which leverages phone calls to guide victims through real-time credential capture and MFA interception, will persist as a repeatable technique for identity compromise and ransomware operations. Threat actors are growing more aware that identity processes (i.e. password resets, MFA resets, helpdesk exceptions) are often less hardened than the technical perimeter, making vishing an increasingly favourable method to achieve persistent access in a short time frame. Vishing poses a significant threat, as one compromised activity can facilitate unauthorized access across an enterprise's SaaS/cloud environment. Even more, vishing is not a technique that was previously explored in similar depth as email phishing; hence, organizations are less resilient to it. Such cases have been widely used during red team operations and the success rate is quite high. Employees are not trained enough to deal with such instances, especially when "call the person" has been a recommendation to mitigate this risk. It is natural for someone who receives a phone call to lower their guard if they have previously trained to use phones as a protection mechanism.
- Vishing has become a key initial access vector in Scattered Lapsus\$ Hunters-branded activity, recently linked to the SSO compromise of identity providers, including Okta, Microsoft Entra, and Google²⁸. Vishing is likely to expand as other extortion and ransomware threat groups attempt to piggyback off SLH's success. Social engineering campaigns in Q2 and beyond will increasingly leverage vishing (especially helpdesk impersonation) for initial access to facilitate lateral movement from credential capture to MFA enrolment, to SaaS session exploitation. This is also echoed by CREST, which, in their 2025 thematic product, called out helpdesk authentication and identity verification weaknesses as an exploitable control gap observed during threat-led penetration testing.
- AI-powered vishing is also expected to rise, by threat actors across the spectrum of sophistication. While targeted deepfakes are already widely observed²⁹, Reversesec's clients have expressed concerns around mass, indiscriminate "agentic vishing" attacks, that would amplify attacker' capabilities beyond physical or linguistic barriers.

²⁷ <https://blog.barracuda.com/2026/01/07/threat-spotlight-phishing-kits-evolved-2025>

²⁸ <https://www.okta.com/blog/threat-intelligence/phishing-kits-adapt-to-the-script-of-callers/>

²⁹ <https://www.resemble.ai/deepfake-database/>

Exercises focused on the latter scenario revealed that inexpensive, publicly available commercial platforms can be (ab)used to rapidly deploy such models, achieving success rates sufficient to warrant large scale escalation.

ClickFix

- As an effective and scalable initial access vector, ClickFix/fake CAPTCHA themes will likely evolve during 2026 but will otherwise remain a consistent TTP due to its ability to target multiple operating systems and bypass defenses by shifting execution onto the user.
- As observed with the CrashFix variant in January 2026³⁰, Q2 will likely see more ClickFix campaigns manufacturing real problems to justify the “fix”, for example, by intentionally crashing the browser, then baiting users into running malicious commands.
- Shortly after CrashFix, in February 2026, a new DNS-based ClickFix variant was observed delivering the emerging ModeloRAT via DNS lookup³¹, demonstrating the continuous and rapid evolution of ClickFix-style techniques that can be expected in Q2 and beyond.

Software supply chain compromise

- The compromise of third parties (software, cloud suppliers, operating systems) through open-source supply chain attacks, especially by breaching third-party software and cloud vendors that leverage AI-enabled tools, will remain a significant threat throughout 2026. AI integrations into third-party services inherit traditional supply chain risk as extra services process sensitive data. This also exposes organizations to information integrity risk (e.g., model tampering and data poisoning).
- Following on a trend that has been building for years and became prevalent in Q4 2025, threat actors increasingly abuse implicit trust embedded in developer tooling, identity, and software distribution to achieve low-noise, high-leverage downstream access, rather than exploiting new vulnerabilities or perimeter access.
- Aforementioned APT operations targeting developers and trusted workflows (including the North Korea-linked Contagious Interview campaign) reinforce that threat actors leverage social engineering to introduce malicious code—an access path that will naturally extend to AI systems wherever AI development and deployment share the same identities, repositories, and build infrastructure.
- Early software supply chain activity focused on hosting malicious packages in non-standard or lesser-known repositories. Q2 2026 will likely see increased targeting of packages hosted in repositories such as Node Package Manager (NPM) and PyPI, Open VSX and other developer tooling platforms. In late 2025, the first wormable software supply chain attack was observed, targeting developer identities and attempting to self-propagate across projects and dependencies. Expect more of these attacks compromising trusted contributors and the distribution of malicious code downstream at scale. This

³⁰ <https://www.huntress.com/blog/malicious-browser-extention-crashfix-kongtuke>

³¹ <https://www.bleepingcomputer.com/news/security/new-clickfix-attack-abuses-nslookup-to-retrieve-powershell-payload-via-dns/>

will likely be very effective, especially where repository monitoring, contributor verification, and dependency integrity controls are immature. AI and agent-based computing are particularly vulnerable to such attacks.

- Through the various CI/CD-focused exercises and “compromised developer” simulations carried out, Reversec has observed a pattern of low or inexistent detective capabilities in SCM / build environments such as Github, Gitlab and Azure Dev Ops. This “common blind spot” constitutes one of the reasons behind the rise in software supply chain campaigns.



Part 4

Testable real-world scenarios

This section turns the threat intelligence in this report into practical adversary models that can be exercised in real environments. Informed by regulated testing and Reversec exercises, the scenarios link tactical intelligence to attack-chain simulation and support resilience validation against techniques expected in Q2 2026.

They cover identity workflows, privileged access paths, developer tooling, and control planes in hybrid and cloud environments. Techniques are mapped to MITRE IDs where available.

Scenario A

Lazarus Group targets financial institutions through the compromise of core infrastructure in a supply chain incident

This scenario was derived from a regulated threat-led penetration test conducted by Reversec for an unnamed organization in the financial sector. It is inherently relevant to the key themes presented throughout this report. While full scenarios are derived from extensive reconnaissance, threat assessment, and understanding of internal and external infrastructure, this is as close as is possible.

Indicative high-level scenario

Reversec attempted to get initial access to the organization by targeting the password reset process. Passwords could be reset remotely for any employee and was handled by Help Desk. The process, roughly, involved an employee calling the Help Desk and asking them to reset their password.

In a mature organization, Multi-Factor Authentication (MFA) is the minimum standard nowadays; therefore, an attempt to reset that was made. Before the call to the Help Desk was made, external

reconnaissance and user enumeration was conducted. As the red team did not know the exact process of the reset password, it was crucial for the operation to try and reset the credentials for an employee that as much as possible was known for. Once sufficient data was collected, a call to Help Desk was placed. Due to poor identity validation, MFA and password were reset and the red team were able to access a virtual desktop infrastructure (VDI) was gained. Since Reversesec operate to a “non-interruption policy for target users”, the engagement continued using synthetic accounts that the organization had created. That approach caused the minimum interruption for the user; their access was disturbed, as they had to reset their credentials again the next working day, but this did not meaningfully affect their ability to discharge their duties.

Once the red team started operating in the environment, reconnaissance took place, revealing a Jenkins installation instance was exposed in a shared folder accessible to any user in the domain. Inside that directory, all the information to decrypt credentials stored in Jenkins were available. Reversesec decrypted those credentials which provided access to the Jenkins web interface. By using the Groovy plugin, it was possible to execute operating system code on the underlying Jenkins server. Although unfiltered access to the internet was not permitted, Reversesec was able to disguise SSH traffic in order to bypass the traffic inspection which enabled a long term, persistent SSH tunnel to be established. This provided an alternative C2 channel to the more traditional implants which had been previously used.

Additionally, various types of credentials were identified. Most importantly, some AWS keys were found that provided access to very specific and sensitive AWS resources. By performing lateral movement on AWS, Reversesec obtained enough permissions to read secrets related to payment systems. As a Proof-of-Concept, Reversesec demonstrated access to a payment system’s database. Due to the nature of the data stored in that database, Reversesec did not proceed further because the objective was achieved.

Scenario B

North Korean IT workers hired in a financial services organization's engineering team, pursue exfiltration of IP, and brokering of long-term access

A long-standing DPRK operation for funding the regime's strategic goals has involved remote IT workers hired under false pretenses in the technology engineering or software development teams of firms across various verticals³². This originally focused on US-based organizations, later expanding to targets globally³³.

Although these operations usually avoid intrusive activities with an aim solely to maintain this position for siphoning of salaries, the hypothesis is presented of the actors receiving new instructions, and adjusting their targeting to pursue higher risk/reward goals such as:

- Access and exfiltration of internal confidential data – including Intellectual Property (IP) such as source code, or sensitive PII – in order to sell it to other parties.
- Backdooring of internal assets with an aim to establish and maintain long-term access to the IT environment.
- Obtaining or creating credentials or other means of shareable remote access into the on-premise or cloud infrastructure, that could be sold to initial access brokers.

Indicative high-level scenario

Reversesec simulated this scenario by performing an “assumed-breach” exercise, from the starting point of a hired remote IT worker with typical developer access, in order to assess the impact of such a scenario materializing. In detail, new employee personas were created, each granted a Virtual Desktop Instance (VDI) as was common for this organizational role profile. Crucially, this profile also involved access to the organization's Source Code Management (SCM) systems, with contribution permissions to assigned projects – as would be typical for a standard developer.

Initially the attack started by performing reconnaissance of the internal environment, searching through file shares and gathering knowledge from internal document stores (T1083, T1213).

With the understanding obtained, the attackers managed to create exfiltration channels leveraging proxy exemptions. These channels were then used to both extract information as well as to infiltrate tooling into the environment (T1105).

³² <https://www.microsoft.com/en-us/security/blog/2025/06/30/jasper-sleet-north-korean-remote-it-workers-evolving-tactics-to-infiltrate-organizations/>

³³ <https://www.okta.com/blog/threat-intelligence/north-korea-s-it-workers-expand-beyond-us-big-tech/>

Reconnaissance also focused on the SCM system, where organization secrets were extracted through weaknesses in Continuous Integration / Continuous Deployment (CI/CD) pipelines configured (T1087). These secrets were subsequently used to perform additional reconnaissance of the cloud identity plane, as well as backdooring of software packages in internal registries and binary distribution platforms (T1525, T1195.002).

Within said CI/CD environment, the ephemeral execution instances (runners) were exempt from the standard security monitoring stack including endpoint detection and response (EDR) tooling. This allowed establishment of Command & Control (C2) channels that were not directly connected to the initial access accounts (T1071).

Exfiltration of larger scale could then be carried out from these runners, allowing all source code repositories to be extracted to attacker-controlled infrastructure. Analysis of the exfiltrated source code repositories for credentials, yielded several valid instances (T1552). This included an access token with administrative access to further repositories, which was iteratively leveraged to configure a backdoor through deploy key. This successfully granted persistent access to the repositories from the public Internet, bypassing enterprise access controls – access that could then be sold to an initial access broker.

Other credential hunting activities carried out from the VDIs discovered privileged credentials in a file share (T1552.001). The corresponding account had control over users of an internal system, allowing lateral movement to their VDIs (T1021). This enabled hijacking of an active session to the internal platform, which would then be leveraged to download and exfiltrate the sensitive data through previously established channels (T1041).

Scenario C

Scattered Spider ransomware operation combining helpdesk vishing, VMWare vCenter exploitation for large-scale disruption

Scattered Spider – also tracked as UNC3944, Oktapus, Octo Tempest, and Storm-0875 among others – is a financially motivated cybercriminal collective that remained active and highly effective in 2025. Unlike many ransomware operators that rely primarily on exploit chains or malware delivery, the group focuses heavily on social engineering and credential compromise to obtain initial access to enterprise environments. Their operations have targeted sectors including telecommunications, retail, aviation, and financial services, often resulting in large-scale data theft or ransomware deployment.

A defining characteristic of Scattered Spider operations is the systematic abuse of human trust within enterprise identity workflows. Combined with native-level proficiency in English, as well as an understanding of the ways of working of global enterprises in the western world, their operations frequently succeed in obtaining initial access through voice-based phishing (vishing) attacks³⁴ against support desks through impersonation of legitimate employees. By convincingly interacting with IT helpdesks, operators manipulate staff into resetting passwords, enrolling new MFA devices, or granting account recovery access. These identity-focused techniques allow the group to bypass many traditional security controls while maintaining a low malware footprint and leveraging legitimate administrative tools already present within the environment.

Following successful identity compromise, Scattered Spider typically escalates privileges and pivots across identity providers, cloud services, and on-premises infrastructure to achieve broad administrative access. Recent campaigns demonstrate a shift toward infrastructure-level compromise, particularly targeting virtualization management platforms such as VMware vCenter³⁵ within vSphere environments. By gaining control of centralized virtualization management, the attack can pivot onto the ESXi hypervisors, providing below-OS level access to critical systems such as Active Directory domain controllers as well as business critical production workloads. It is through this access – traditionally unmonitored by sophisticated security solutions such as EDR – that credential databases can be skimmed and exfiltrated, and encryption of disks performed, in what is known as a double-extortion attack.

³⁴ <https://www.crowdstrike.com/en-us/blog/crowdstrike-services-observes-scattered-spider-escalate-attacks/>

³⁵ <https://cloud.google.com/blog/topics/threat-intelligence/defending-vsphere-from-unc3944>

Indicative high-level scenario

Following the escalation of Scattered Spider’s disruptive operations in 2025 attracting industry attention, Reversesec was engaged for numerous exercises of different formats. These exercises all had in common a motivation to proactively assess resilience driven by Scattered Spider’s campaigns, and an intelligence-led approach for the selection of simulated TTPs against applicable attack surfaces. Two of these exercises present particularly insightful examples:

- A detection-focused Purple Team against the VMware vSphere estate of a global bank.
- A social engineering exercise against the helpdesk of a global financial services provider.

By combining the activities carried out in these two instances, an indicative scenario can be formed for simulating a realistic Scattered Spider campaign end-to-end. The details of such a scenario are laid out below, with TTPs involved organized across several different phases. Organizations seeking to reproduce this plan end-to-end are advised to design “leg-ups” or pre-agreed provisions that could be used by the red team to progress along each stage, should no viable attack path be found.

Phase 1: Vishing for initial access

In this first stage, Scattered Spider’s helpdesk campaigns can be simulated through a series of social engineering calls against the service / support desk phone numbers advertised publicly, with the aim of getting access as a legitimate employee through a password and/or MFA reset.

Various pretexts can be explored by operators such as:

- New hires calling to get onboarded.
- Users without a cloud workspace, calling to be assigned one.
- Contractors from known service suppliers or vendors.
- Device loss or theft, combined with distress conditions.

Reversesec advises that caller ID spoofing techniques are employed as well, to improve perceived legitimacy prior to the conversation even commencing, while simultaneously testing for potential procedure weaknesses due to trust on spoofable indicators such as calling number.

Phase 2: Identity privilege escalation and pivot to vCenter management interface

If a foothold is achieved, or a leg-up is used to provide initial access, the privilege escalation stage of Scattered Spider’s

operations can be simulated through internal activities. The group seeks elevation of privileges as a means to pivot to the virtualized infrastructure control plane, which is typically integrated with identity platforms such as Active Directory. Red teams could perform a subset of common Active Directory attacks, or credential hunting activities, which align with threat intelligence from documented incidents.

Once administrative access is achieved, manipulation of the relevant security groups would follow (T1098.007), that would grant Single Sign-On (SSO) access to the web-based VMWare vCenter control plane. Alternatively, if a user with such access is successfully compromised (or assumed compromise) irregular logins and UI activities like the following could stress any behavioral detections present:

- Creation of rogue users, adding them to groups, modification of user privileges (T1136.001, T1098.007).
- VM related operations such as creating rogue VMs, taking memory dumps and disk snapshots (T1005).
- Cloning, copying, and hijacking of disk images (VMDKs) which could allow subsequent credential extraction (T1552).
- Attempts to compromise VM guests via boot order manipulation, backdooring of ISO media (T1542), reconfiguration of virtual network interfaces (vNICs) and launching of console sessions.
- Data exfiltration techniques involving Copy/Paste utilities, VMDK downloads from the UI.
- Obtaining shell access to the vCenter appliance through an SSH session (T1021.004).
- And finally, disabling or modifying Lockdown mode controls for ESXi hosts, to facilitate progression to the next stage (T1562).

Phase 3: ESXi execution and persistence

The next phase of the scenario would model Scattered Spider's final attack positioning step, the lateral movement to the hypervisor environment, through SSH access to the ESXi shell (T1021.004) or web access to the ESXi Host Client UI.

Indicative actions for this phase could be executed using three different means:

- **Python** execution – Through the native Python interpreter present on ESXi hosts (T1059.006).
- **Pre-compiled binaries** (C/C++) – Statically compiled binaries for the target architecture, could be used to interface with system APIs allowing custom code to be executed (T1106).
- **"Living Off The Land" utilities** – Finally, the native utilities already present on the ESXi shell could be used (T1059.012).

These execution methods would allow simulation of the following TTPs. It is noted that these do not strictly adhere to Scattered

Spider intelligence, but incorporate techniques observed by other threat actors such as UNC5221³⁶, UNC3886³⁷, LockBit³⁸, and FireAnt³⁹:

- Downloading (T1105) and running pre-compiled binaries to establish a network tunnel with other components, or achieve Command and Control (C2) (T1572).
- Disabling system security controls such as “execInstalledOnly” and creating rogue VMs outside of standard process to operate unmonitored (T1562).
- Backdooring “rc.local” scripts to execute on boot, achieving persistence (T1037.004).

Suspending VMs on the host in preparation for ransomware deployment (T1489).

Phase 4: NTDS exfiltration and VM encryption

The final phase of the scenario would involve simulation of the actions on objectives through the ESXi shell.

- Exfiltration of Active Directory database data (NTDS.dit) from domain controller VMs (T1003.003, T1041).
- Exfiltration of VM snapshots and disk images containing sensitive enterprise data (T1005).
- Mass encryption of VMDK files to render hosted virtual machines unavailable (T1486).

36 <https://cloud.google.com/blog/topics/threat-intelligence/brickstorm-espionage-campaign>

37 <https://cloud.google.com/blog/topics/threat-intelligence/uncovering-unc3886-espionage-operations>

38 https://www.trendmicro.com/en_us/research/22/a/analysis-and-impact-of-lockbit-ransomwares-first-linux-and-vmware-esxi-variant.html

39 <https://www.sygnia.co/blog/fire-ant-a-deep-dive-into-hypervisor-level-espionage/>

Get in touch!

Book a threat review session to turn these findings into action.

Prioritize your security investments or plan your next Red Team exercise.

[BOOK NOW](#)





About ReverseSec

ReverseSec, a new name in cybersecurity consulting, helps organizations worldwide tackle their most complex cybersecurity challenges.

With a focus on continuous research, practical solutions and knowledge sharing, ReverseSec's findings provide the rationale behind informed security decisions.

With over 30 years of experience, ReverseSec brings together the expertise of renowned companies MWR Infosecurity, F-Secure, WithSecure, Digital Assurance, nSense, and Inverse Path.

ReverseSec.com

REVERSESEC