Ontinue

1 H 2025

Threat Intelligence Report

Executive Summary

The first half of 2025 confirmed that cyber threats remain relentless, with ransomware, phishing, and state-aligned activity continuing to evolve in sophistication and scale.

Ransomware: Still a Top Threat

Despite a 35% drop in reported ransomware payments in 2024 ($1.25B -> $813M), attacks remain frequent and damaging. There were 4,071 claimed ransomware breaches in H1 2025 across 109 countries, driven by 90 active groups, led by CLOP, AKIRA, and QILIN. Services, manufacturing, IT/communications, and retail/wholesale sectors were most targeted.

Phishing and Identity Attacks

Phishing-as-a-Service matured with Tycoon 2FA, responsible for ~65% of PhaaS-based credential attacks observed by Ontinue. This AiTM platform bypasses MFA to compromise Microsoft 365 and Gmail accounts. Weaponized SVG phishing campaigns rose 40%, using embedded scripts to evade email filters and exploit user trust.

Cloud Persistence and the Red Team Gap

Analysis in the Ontinue Cyber Defense Center revealed a widening gap between red team exercises and live adversary behavior in Azure environments. While simulated campaigns highlight resilience under controlled conditions, real-world attackers employ persistence methods without rules of engagement, extending dwell times and enabling monetization.

- Over 70% of phishing attachments bypassing secure email gateways in H1 2025 were non-traditional formats such as SVG or IMG.
- Roughly 20% of incidents investigated involved refresh token replay, allowing adversaries to persist after password resets and bypass MFA.
- Nearly 40% of Azure persistence cases involved multiple, layered methods for redundancy.
- Successful intrusions often included tampering with diagnostic settings or conditional access policies, extending dwell time to a median of 21 days.

Malware Trends

The Lumma C2 infostealer, linked to at least 1.7M credential thefts, suffered a major takedown in May 2025, with 2,500 domains seized. Despite disruption, resilient infrastructure suggests continued risk.

A notable resurgence was also observed in USB-delivered malware, with a 27% increase compared to H2 2024. Though far from new, these attacks remain effective, leveraging removable media to bypass network defenses and introduce malware directly onto endpoints.

Advanced Threat Actors

- Scattered Spider – Blends social engineering with cloud exploitation, often via trusted third-party vendors.
- Predatory Sparrow – Pro-Israeli actor targeting Iranian financial and industrial assets.
- Void Blizzard – Pro-Russian espionage against NATO-aligned critical infrastructure.
- Lazarus Group – North Korean state actors responsible for a $1.5B Bybit crypto heist.

Third-Party Risk

Vendor-related breaches doubled YoY, now implicated in ~30% of incidents. Weak security in external partners facilitated attacks on M&S and Adidas, underscoring the need for robust vendor risk management.

TLDR

Criminal and state-aligned actors are adapting faster than ever, targeting weak links across the technical, human, and third-party spectrum. Only layered, adaptive defenses backed by current threat intelligence will enable organizations to reduce risk in the months ahead.

Table of Contents

CHAPTER 1

Ransomware

Ransomware Still a Long-Term Threat

Evolving Tactics of Ransomware Operators

Ransomware activity remained a significant global threat in the first half of 2025. While reported ransom payments declined 35% in 2024, this trend appears linked to stronger resistance to payment and law enforcement actions, rather than a reduction in attacks.

Data shows 4,071 claimed ransomware breaches between January and June 2025, involving 90 distinct groups and impacting organizations in 109 countries. The most active groups included CLOP, AKIRA, and QILIN. Services, manufacturing, IT/communications, and retail/wholesale were among the most frequently targeted sectors.

Notable incidents included disruptions at major UK retailers and ongoing recovery efforts from prior high-profile attacks, illustrating ransomware's continued operational and impact across industries.

For example, the British Library, one of the world's biggest and most important institutions of its type, was attacked by criminals affiliated with the Rhysida ransomware group in late 2023 and hit with a blackmail GDN demand for close to than £600,000 (then about $450,000), which it didn't pay.

https://www.theguardian.com/commentisfree/2024/feb/06/hacker-british-library-cybersecurity-cybercrime-uk

Library services were entirely disrupted at first, and reverted to 1970s-style pen-and-paper access:

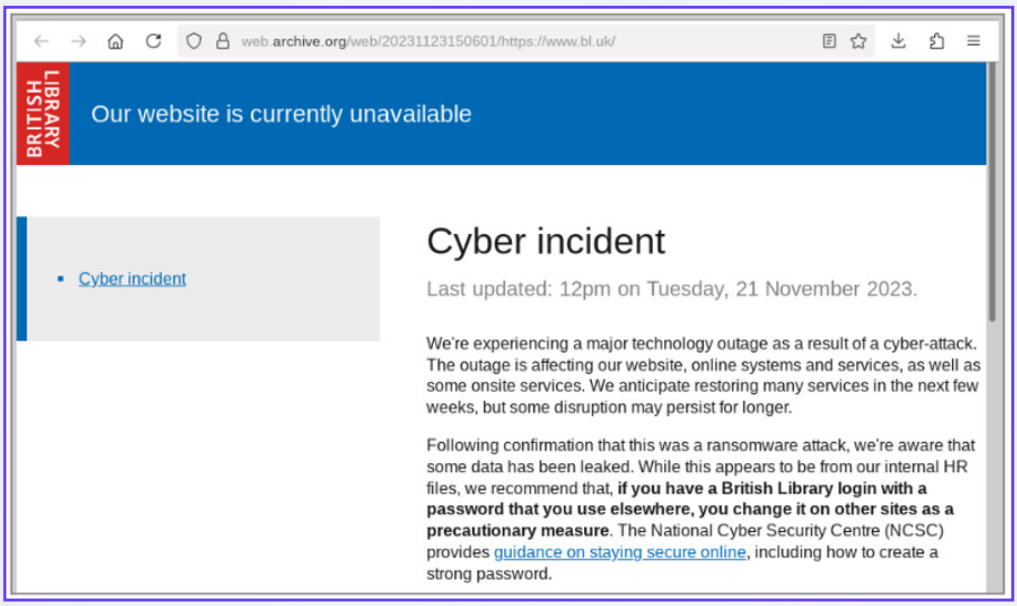

Partial service was eventually resumed and continued throughout 2024 and H1 2025, with the library's official home page noting that "[t]his is a temporary website, with limited content" right up to 2025-07-06:

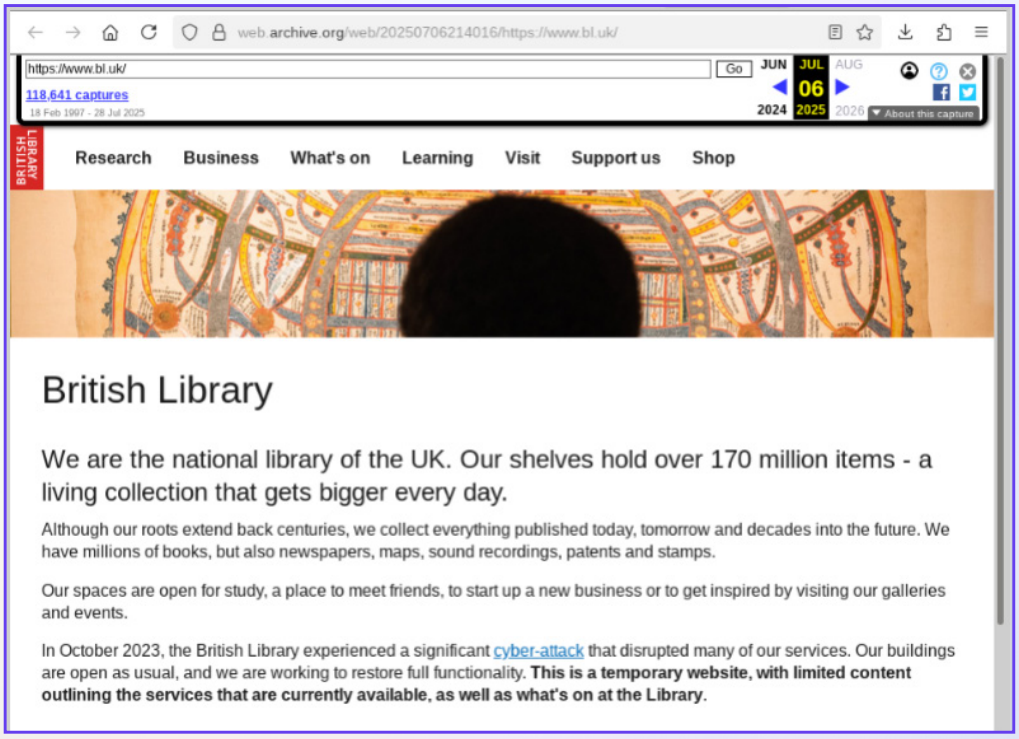

Full service was restored only on July 7, 2025:

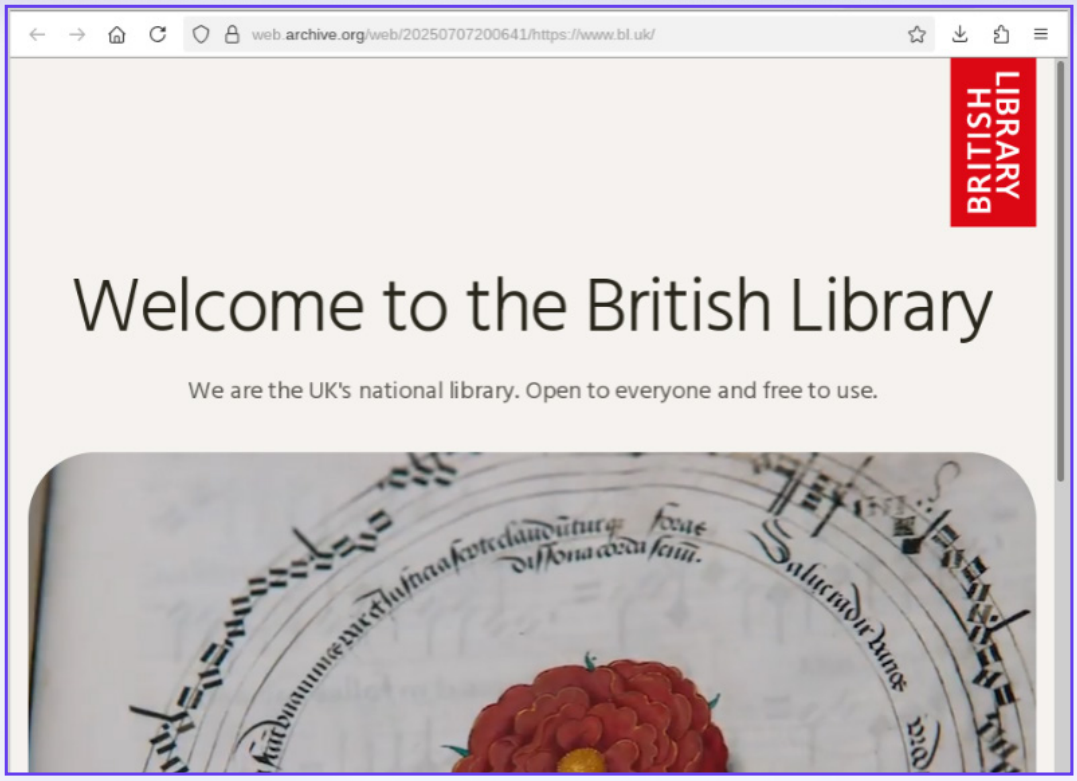

Affiliate Networks are Fragmenting but not Vanishing

One common and unfortunate side-effect of ransomware takedowns is that the affiliates of the now-blocked "service" are unlikely to get caught up directly in the takedown.

The affiliate-style system used by so-called RaaS (ransomware-as-a-service) groups relies not on written contracts but on pseudo-anonymous, hard-to-trace allegiances agreed via the dark web.

As a result, those affiliates often end up shopping around for other groups to join, bringing criminal knowledge with them: practical experience from previous network intrusions, ransomware playbook information from the now-defunct group, and decryption keys and data stolen from victims who haven't yet paid up. For example, the LockBit takedown in 2024 was followed by what seems to have been a scramble by some of the group's core members, or at least some of its affiliates with access to the core servers and databases, to rebuild a new-look LockBit, dubbed simply LockBit 4.0.

This resurrected crime group was ambitiously and arrogantly promoted on the dark web in December 2024, with a claimed launch date of early February 2025. This was followed by a "recruitment campaign" offering affiliate status to anyone willing to pony up $777 in Bitcoin, and by the release of yet another ransomware variant, unsurprisingly called LockBit 5.0.

However, by May 2025, LockBit's criminal service empire had been broken into TREL again - this time not by law enforcement but by an unknown attacker identifying themselves with the tagline xoxo from Prague.

https://www.trellix.com/blogs/research/inside-the-lockbits-admin-panel-leak-affiliates-victims-and-millions-in-crypto/

The attacker dumped a SQL database covering LockBit operations from December 2024 to late April 2025, including information on at least 156 victims, of whom more than 100 had apparently entered into negotiations with their attackers.

The existence of this dumped data, along with its analysis, suggests that both operational security and discipline among ransomware affiliates are very poor indeed, which creates a very strong argument against paying off attackers.

For example, the data stolen from LockBit apparently revealed:

- Information about affiliates and victims, including details of negotiations conducted and cryptocurrency wallets used.
- Uncertainty about whether individual affiliates or the core group members controlled the needed decryption keys, suggesting that paying up would not help to unscramble ruined data.
- Uncertainty in who had access to already-stolen data, suggesting that many affiliates belong to multiple RaaS groups and shuttle data between them, which implies that paying up gives little "protection" against the subsequent disclosure of that data, or against a second extortion attempt.

https://solcyber.com/how-far-can-you-trust-a-ransomware-criminal/

The Challenge of Ransomware Payments in 2025

The decision of whether to pay a ransomware demand remains one of the most complex issues organizations face following an attack. In 2025, the stakes have grown higher as most ransomware incidents now combine large-scale data theft with encryption of critical systems, increasing both operational and reputational pressure on victims.

Originally, ransomware relied on encrypting files in place as a faster alternative to data exfiltration, immediately disrupting business operations while holding recovery hostage to a decryption key. Today, attackers often do both—scrambling data to halt operations while exfiltrating sensitive "trophy" data to use as additional leverage.

Paying for a decryption key is at least a transaction with a measurable outcome, even if the tool underperforms. Paying for a promise to delete stolen data is different—it's an unverifiable assurance that the information will be erased and not resold or leaked. This risk was underscored in April 2025, when the LockBit group itself suffered a breach that exposed affiliate and victim information, demonstrating that even ransomware operators are not immune to compromise. Regulatory approaches are beginning to shift in response. In early 2025, Australia introduced rules requiring organizations to report any ransomware-related payment—or provision of any benefit—to authorities within 72 hours. The regulation covers all forms of blackmail related to a cybersecurity incident, regardless of payment method or whether a third party negotiated on the victim's behalf. The UK has announced plans for similar requirements, alongside a proposed ban on ransom payments by public sector entities.

While some experts advocate for an outright ban on all ransomware payments, most governments remain cautious, concerned that prohibiting payment entirely could push transactions underground, reduce incident reporting, and impede investigations. Current trends indicate a gradual move toward greater transparency and regulation rather than blanket prohibition.

Attack Metrics for H1 2025

The apparent fall in ransomware payments during 2024 does not seem to have translated into an obvious decline in the overall risk of, and threat from, ransomware-style attacks during H1 2025.

Measuring the number of attacks is a tricky business, not least because there are few reliable sources of ransomware reports.

As mentioned above, Australia now requires most ransomware payments (even if made indirectly via a proxy or an out-of-country "negotiator", or if made in some non-monetary form) to be reported within 72 hours, unlike in most countries. Even in Australia, the new reporting rules don't generate any reliable data about attacks where no payment was made.

And many countries, even those with mandatory data breach reporting regulation, seem to tolerate network intrusions going unreported when the victim pays hush money to the intruders to keep quiet about the incident.

Paying the blackmail may greatly reduce the risk of the stolen information being abused in future, but it doesn't in any sense "unsteal" or "de-breach" that data. We have therefore used data collected from ransomware breach claims made on the dark web by threat actors themselves, which gives some fascinating insights into the nature and scale of the problem, albeit with two important caveats:

False positives.
The criminals themselves are neither objective nor reliable witnesses. Intrusions where no data was stolen may be reported as if they had succeeded; breaches might be exaggerated in size; and incidents may be reported more than once in ways that can't reliably be deduced from the data. Affiliates may attempt to "advertise" the same breach in different ways on multiple online forums, hoping to attract multiple buyers for the same data.

False negatives.
Breaches where victims paid up promptly may deliberately go unmentioned by the attackers, as part of their extortion "promises" to keep attacks quiet. Attacks where businesses were disrupted but no trophy data was exfiltrated may never be claimed, because the criminals lack sufficient "proof" to support their dark web bragging.

Two of the most notable events include the Marks & Spencer and the Co-op breaches in the UK. These attacks apparently relied on DragonForce ransomware code and together affected about 125,000 employees and at least 6.5 million customers.

Nevertheless, the breach claims make dramatic reading, with:

- **4071 claimed breaches** in the six-month data set for H1 2025 (2025-01-01 to 2025-06-30 inclusive).
- **90 differently-named ransomware groups** involved in these breaches.
- **At least 109 countries** were affected.
- **Organizations of all sizes are impacted**. Sole proprietors, SMEs and Enterprises are all attacked.

(Not all claimed breaches could be connected to a specific country or organization. These were excluded when computing the numbers below.)

Top 12 Ransomware Groups by Breaches Claimed

With 90 groups in the data, it's clear that ransomware criminality is both prevalent and popular.

But the top groups in the list are especially active (or at least especially inclined to brag about their crimes).

The top seven groups in the list all averaged more than one attack per day, with CLOP and AKIRA claiming more than two victims per day, and that's just for attacks that the groups chose to announce online.

This is strong evidence of the "force multiplier" effect of the affiliate model used by ransomware groups to draw new criminals into the fold:

Name	Ratio	Count
CLOP	10.1%	411
AKIRA	9.4%	382
QILIN	8.4%	344
RANSOMHUB	5.8%	236
PLAY	5.3%	214
SAFEPLAY	4.7%	191
LYNX	4.5%	183
BABUK 2.0	3.7%	150
INC RANSOM	3.3%	133
MEDUSA	2.6%	107
FOG	2.2%	91
DRAGONFORCE	2.0%	82

Top 12 Countries by Organization Attacked (Where known)

Unsurprisingly, perhaps, the US topped the list, though that figure might be skewed due to attacks being counted against just one country. (Many multinationals list their headquarters as the US, even though attacks against such organizations may affect staff and customers in many countries)

Interestingly, of the top 12 most populous countries in the world, only Russia doesn't appear anywhere in the 109-country list.

Although we know from recent reports about the LockBit group that Russian companies have indeed been targeted this year, many ransomware groups operate out of Russia, and deliberately avoid targeting Russian businesses to since it's an increased risk, so many Threat Actors target foreign victims. (Russia is one of several countries in the world that prohibits its own citizens from being extradited to face trial abroad.)

Name	Ratio	Count
US	54.5%	1925
CANADA	6%	213
GERMANY	4%	142
UK	3.7%	129
ITALY	2.4%	86
SPAIN	2%	70
FRANCE	2%	69
BRAZIL	1.9%	68
AUSTRALIA	1.8%	65
INDIA	1.5%	52
TAIWAN	1.1%	39
SINGAPORE	1%	36

Organizations Attacked by Sector (Where known)

As in our H2 2024 report, the manufacturing and services industries were the most attacked.

Worryingly, perhaps, the IT/Communications sector was also hit at an above-average level.

Clearly, however, no sector is immune.

Name	Ratio	Count
CONSTRUCTION	8.2%	288
EDUCATION	3.4%	118
FINANCE / INVEST	9.5%	333
GOV / SOCIETY	5.8%	202
HEALTH / PHARMA	7.8%	271
IT / COMMS	10.8%	378
MANUFACTURING	12.9%	451
POWER / FUEL	3.8%	132
PRIMARY INDUSTRY	2.8%	99
RETAIL / WSALE	9.6%	335
SERVICES	16.2%	567
TRANSPORT	9.1%	317

Victims by Number of Employees (Where known)

If we take the average size of each band below (assuming that the average "over 1000" organization has 4000 staff), these numbers represent more than 2 million employees whose personal data may have been breached and stolen by cybercriminals to use as blackmail leverage.

This alone is a serious warning of the potential cybersecurity risk of any ransomware attack, even if business operations were not directly disrupted and no ransom was paid.

No business can rely on being too small to be of interest. Small companies are often an entry point for new gangs attempting to climb the ladder.

The Co-op breach stands out among the 2025 H1 high-profile data incidents, which exposed the personal data of 6.5 million customers and up to about 60,000 employees.

Size	Count
Up to 50	1,112
51 to 200	1,012
201 to 1000	745
Over 1000	397

CHAPTER 2

Threat Spotlights

Tycoon 2FA

Phishing isn't just an inbox annoyance anymore it's a highly industrialized, scalable threat powered by sophisticated toolkits and clever social engineering. One of the standout platforms fueling this evolution is Tycoon 2FA, a Phishing-as-a-Service (PhaaS) operation that's been active since 2023 and surged in popularity in early 2025. While previously noted in the wild, Ontinue observed a significant spike in its activity during the first half of 2025, suggesting broader adoption among threat actors and potential campaign coordination.

Ontinue CDC (Cyber Defense Center), found that Tycoon2FA was responsible for approximately 65% of all phishing-as-a-service (PhaaS)-based credential attacks in the first half of 2025. This dominance not only demonstrates the kit's effectiveness but also highlights the extent to which it has become the preferred choice for adversaries operating at scale.

What sets Tycoon 2FA apart is its focus on Adversary-in-the-Middle (AiTM) phishing, particularly targeting Microsoft 365 and Gmail users. This isn't your average spoofed login page. Tycoon 2FA combines advanced anti-detection features, multi-layered obfuscation, and 2FA interception capabilities making it a go-to platform for cybercriminals seeking high success rates.

How the Attack Works

It usually starts with a well-crafted spear phishing email, often coming from a compromised account or a spoofed trusted contact. These messages contain links or QR codes leading to malicious pages designed to bypass automated detection systems. Victims are initially shown a CAPTCHA or Cloudflare Turnstile, helping weed out bots and sandboxes before progressing to the actual phishing payload.

Once on the malicious site, victims are subjected to a series of security checks designed to frustrate researchers and avoid analysis. The phishing kit disables browser inspection tools (like F12 or right-click) and looks for signs it's being observed such as virtual machines, security tools like Burp Suite, or browser automation frameworks like Selenium.

Only if the environment passes these checks will the user be redirected to a convincing spoof of the Microsoft 365 login page. But it doesn't stop there: after entering their credentials, victims are shown a fake 2FA prompt. Any tokens entered here are immediately harvested and exfiltrated along with session cookies to give attackers full access, bypassing MFA entirely. Finally, the user is redirected to a legitimate Microsoft site to minimize suspicion.

The Infrastructure Behind It

Tycoon 2FA isn't just a clever front-end trick. Its infrastructure is agile and constantly shifting. Ontinue's ATO team observed phishing domains registered using typo-squatting tactics (e.g., microsOft-login[.]com) with short life spans and protection via Cloudflare or other reverse proxies. These campaigns are typically delivered through hijacked email accounts or abused cloud services, further boosting their credibility.

The takeaway? Tycoon 2FA represents the future of phishing modular, evasive, and easily accessible to less technical attackers through PhaaS models. Organizations need to move beyond simple email filters and start defending against entire ecosystems of tools that enable these campaigns.

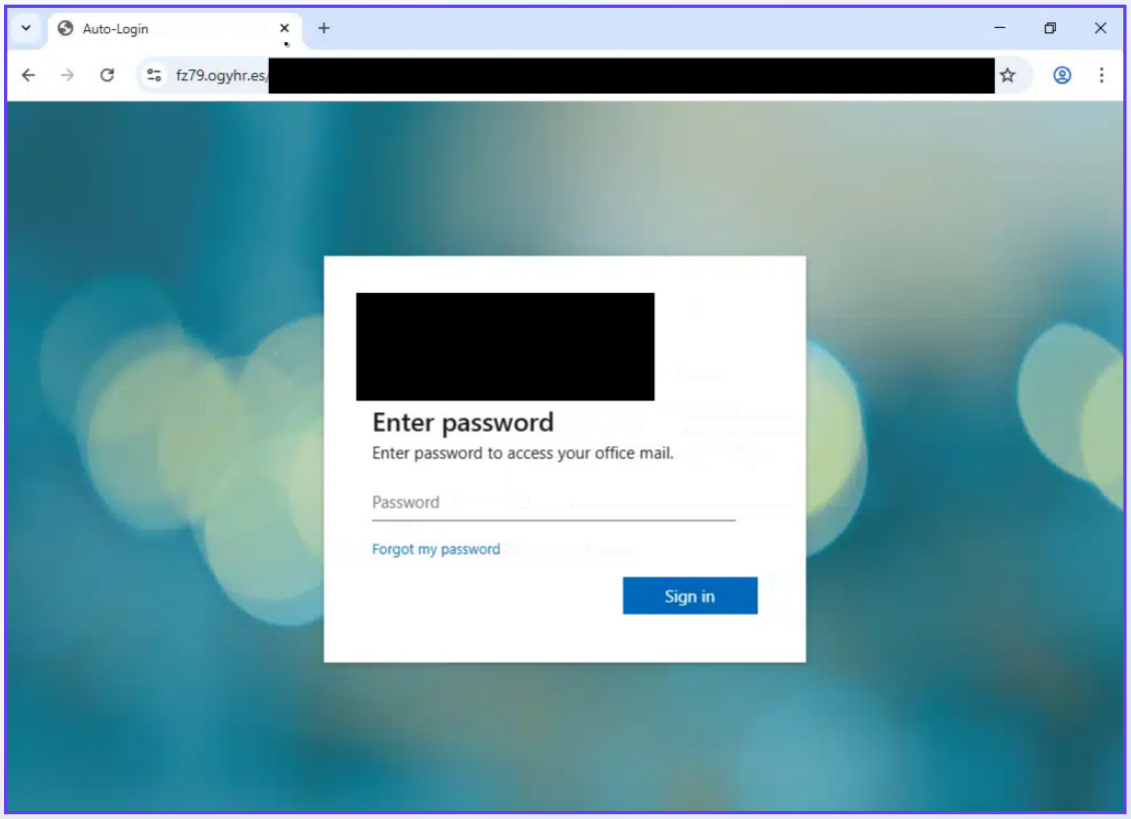

Weaponized SVGs

Ontinue has observed an increasing use of weaponized SVG files in phishing attacks. While this technique has been documented publicly as early as 2016, it continues to be actively leveraged by threat actors today. These files often contain embedded and sometimes obfuscated JavaScript, which is used to perform redirects or load malicious content. Such attacks rely on social engineering tactics and exploit the HTML5/SVG rendering capabilities of modern browsers to bypass traditional email security filters that typically treat svg files as harmless image formats.

Why It's Effective

- SVG is XML-based and scriptable, unlike most other image formats.
- Email security tools historically didn't inspect SVG internals (they were often whitelisted).
- SVGs can be previewed natively by email clients or browsers, triggering execution with minimal user interaction.

User Behavior

- Why users trust it, it's just an image.
- Lure technique: fake invoices, HR forms, Delivery notifications or even
- Behavioral red flags to watch for (e.g., previewing unknown SVGs in webmail

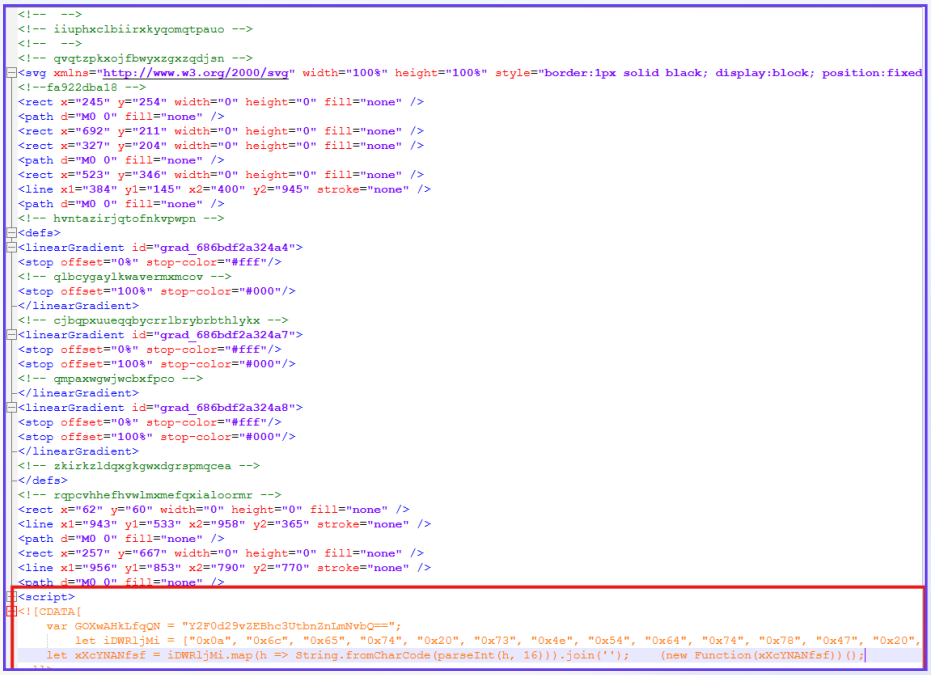

Observed Trends

Ontinue has observed a 40% surge since the beginning of the year in weaponized SVG files being leveraged in phishing campaigns. While SVGs have appeared in phishing attacks sporadically in previous years, recent months show a marked increase in their use as an initial access vector.

This resurgence has coincided with the increased deployment of the Tycoon1FA phishing kit, suggesting that threat actors are adopting SVG-based delivery mechanisms as part of more sophisticated multi-factor phishing operations.

Scattered Spider

There's a new breed of threat actor blending sharp technical skills with even sharper social engineering and Scattered Spider is leading the charge. Also known as Octo Tempest, Starfraud, UNC3944, Scatter Swine, Oktapus and Muddled Libra, this financially motivated group has been active since 2022, carving a name for itself by relentlessly targeting large enterprises across multiple industries.

What makes Scattered Spider especially dangerous isn't just their malware toolkit (though they use stealers like AveMaria, Raccoon, and VIDAR). It's their mastery of impersonation and manipulation. They don't break in by brute force, they talk their way in, posing as IT support staff, convincing employees to install remote access tools like ScreenConnect, TeamViewer, or Pulseway, and even persuading them to approve MFA prompts or hand over authentication codes.

A Shifting, Expanding Target List

Scattered Spider doesn't stay in one lane. They've hit organizations in hospitality, retail, finance, telecom, gaming, crypto, and more. According to the FBI, their latest area of interest is aviation, meaning any company that's part of the broader airline ecosystem, including third-party vendors, could be in their crosshairs.

Once inside, they move quickly and quietly. Their attack methods span from the tactical, like creating new user accounts or modifying MFA tokens, to the highly technical, such as federating rogue identity providers into the victim's SSO infrastructure to escalate privileges and maintain persistent access.

Breaking Down Their Playbook

Here's how their attacks often unfold:

- **Initial Access:** It starts with social engineering. Employees are tricked into installing remote tools or resetting credentials. The group often abuses trusted relationships with third-party IT support vendors to legitimize their deception.
- **Execution & Persistence:** Once inside, they use commercial remote access tools and serverless cloud functions to move undetected. They create new identities, maintain access using valid credentials, and manipulate MFA configurations to avoid losing their foothold.
- **Privilege Escalation & Evasion:** The attackers elevate privileges by linking their own identity providers into the victim's systems. They also spin up new cloud instances and impersonate IT staff to continue harvesting credentials and moving laterally.

- **Credential Theft:** Scattered Spider uses stealer malware, pushes repeated MFA notifications (aka "push bombing"), and scours systems for insecurely stored credentials and keys.
- **Discovery & Lateral Movement:** Their reconnaissance is methodical, they hunt through SharePoint, backups, and browser histories, and even use AWS Systems Manager Inventory to identify potential lateral movement paths.
- **Data Collection & C2:** Data is staged from across the environment into a central location before exfiltration. Communication with command-and-control servers is often managed through legitimate remote tools, helping them blend in with normal traffic.

Why This Group Matters

Scattered Spider represents the convergence of cloud abuse, identity manipulation, and good old-fashioned social engineering. Their campaigns show how attackers don't need zero-days when they can just talk their way in, and why human-layer defenses like awareness training, identity governance, and MFA hardening are just as critical as endpoint protection.

USB Malware and Basic Exposure Risks

While much of the cybersecurity conversation in 2025 centers on advanced threats such as ransomware-as-a-service and AI-enabled attacks, incidents in H1 2025 show that older techniques remain both relevant and effective. Ontinue's Advanced Threat Operations (ATO) team observed a 27% increase in USB-delivered malware compared to H2 2024, underscoring that even long-standing attack methods continue to impact modern enterprises.

This attack method is not new. USB-delivered malware has been in use for decades, but it remains a viable option for adversaries seeking initial access. According to

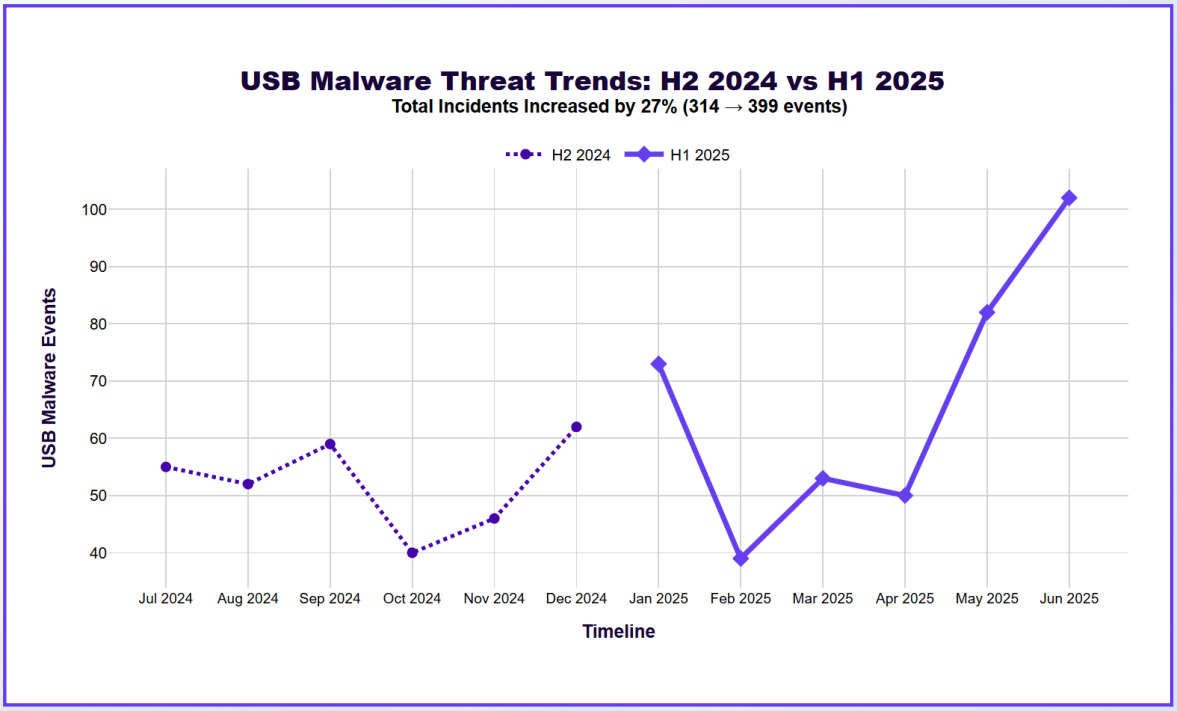

a 2024 Honeywell report, over half of USB-based threats (51%) had the potential to cause major disruption to industrial and enterprise environments. Despite the availability of device control capabilities and advanced endpoint protection, many organizations continue to allow the use of removable media without strong restrictions.

The implications are significant. Malware introduced through a single USB connection can bypass network-level defenses and initiate an infection chain with operational consequences disproportionate to the simplicity of the action. In one recent case, connecting a personal USB drive to a managed workstation triggered a malware infection that required rapid containment to prevent further spread. Such incidents highlight that user actions and policy gaps, rather than sophisticated zero-days, often provide the initial foothold for attackers.

Similar risks arise when devices are exposed through misconfigurations or unauthorized remote access. In several instances, attacker reconnaissance and credential-based access attempts originated not from novel exploitation, but from user-created tunnels or misconfigured systems. These methods demonstrate that attackers are prepared to exploit whichever path requires the least effort, whether through technical sophistication or human error.

Key Defensive Measures:

- Restrict or monitor USB device usage across managed environments.
- Minimize local administrative privileges to reduce misconfiguration risks.
- Continuously monitor for exposed assets and unauthorized tunneling.
- Reinforce user awareness through training and policy enforcement.

These findings emphasize that older, straightforward attack vectors remain effective because they exploit predictable behaviors and overlooked fundamentals. Organizations that focus exclusively on advanced threats risk leaving themselves exposed to low-complexity, high-impact compromises.

Lumma C2 Malware

Lumma C2, one of the most prevalent information-stealing malware families in 2024 and 2025, reported significant disruptions to its operations after a law enforcement takedown in May 2025. This news emerged following the press release from Europol stating that they have seized almost 2,500 domains that are controlled by Lumma MaaS.

Operational Impact

The US Department of Justice along with Microsoft DCU, and European and Japan governments has taken down the 2300 internet domains that are used by LummaC2 actors and their affiliates. This operation has targeted the malware's infrastructure and user panels that caused disruption to the service.

The FBI has identified at least 1.7 million instances of where this malware was used to steal information, making it the most prolific infostealers. As seen from the chats on cybercrime platforms, the seizure has noticeably impacted the reputation of Lumma leading to the decrease of its usage. Despite the disruption, Lumma's developers have swiftly started claiming to be operational and working normally. Excerpts from the Telegram chats of cyber criminals with Lumma's developers posted online show them saying that no one related to Lumma was arrested and "everything has been restored, and we are working normally."

The core strength of Lumma lies in its infrastructure and the service model that provides accessibility that led to a rise in popularity amongst cyber criminals.

They have claimed that the operation done by FBI has not impacted their main server because of the region it's located in. Although they were able to penetrate their server and extract information through credential harvesting and digital footprints, their server is still undisturbed as per their claims.

In another instance, just 2 days after the operation by law enforcement agencies, an automated Telegram message was seen offering a sale of stolen credentials from Lumma. In conclusion, despite the successful seizure of thousands of Lumma domains, malware seems to be prevalent and the demand for infostealers on the dark web market has not come down. Lumma's resilient infrastructure suggests it could resurface potentially under a new name, while maintaining the same tactics and impact.

Tactics, Techniques, and Procedures (TTPs)

Initial Access:

- Phishing emails with malicious attachments or links
- Malvertising on cracked software and adult content websites
- Loader malware used to deploy Lumma as a second-stage payload

Execution:

- Deployed as a packed binary via obfuscated loaders
- Uses native Windows APIs for stealth and to avoid detection

Persistence:

- Minimal to none; Lumma typically operates in-memory or is short-lived
- May configure autoruns or use scheduled tasks if instructed

Data Exfiltration:

- Steals credentials from browsers, crypto wallets, autofill data, and session tokens
- Exfiltration over HTTP(S) to C2 domains or hardcoded IPs
- Uses AES for encryption before transmitting data

Evasion:

- Code obfuscation and dynamic API resolution
- Sandbox and VM detection to evade automated analysis
- Command and Control:
- Communicates with PHP-based C2 panels
- Previously utilized .top, .xyz, and .shop domains for distribution, we are starting to observe newer domain TLD's

CHAPTER 3

Security Testing vs. Real Incidents

Beyond the Simulation: **Aligning Security Testing with Real-World Threats**

Real-world threats differ in important ways from penetration testing and adversary simulation, despite sharing similarities. Penetration tests are simulated attacks, and while they aim to mimic real-world scenarios, they might not always capture the full complexity and adaptability of a malicious actor. While pentests and attack simulations are controlled, scoped exercises designed to evaluate specific aspects of security posture, real-world adversaries are unpredictable, adaptive, and unconstrained by predefined rules of engagement. Threat actors often exploit overlooked systems, social dynamics, or operational weaknesses that simulations may miss. Additionally, some organizations prioritize penetration testing primarily for compliance purposes rather than as a core part of their overall security strategy, which can result in superficial assessments.

This highlights the importance of aligning defensive strategies with real threat intelligence grounded in current tactics, techniques, and procedures (TTPs) observed in the wild, rather than relying solely on test-driven scenarios. Bridging this gap enhances resilience and ensures defenses are relevant to today's evolving threat landscape. Prioritizing defenses and detections based on real security incidents observed within your organization ensures that resources are focused where threats are most active and impactful. Unlike hypothetical scenarios, which can lead to diluted efforts and alert fatigue, data-driven prioritization grounds your security posture in actual adversary behavior. This approach increases detection efficacy, improves incident response, and aligns defenses with the specific tactics, techniques, and procedures (TTPs) targeting your environment, delivering measurable risk reduction where it matters most.

Selecting the primary objective of Security Testing is important in multiple ways. Domain-wide assessments starting from the "classic" compromised laptop scenario will result in potentially missing the strengths of the security program that focuses on preventing initial access. "Inside the network" scenarios overlook the advantages of having endpoint telemetry.

Testing key systems is understandable yet might miss the mark to safeguards assets through lateral compromises. Running regular crown jewels focused penetration tests and using the reports to fix the findings is often narrow-focused - if you have a deep understanding of the strengths and weaknesses of the technical controls of your security environment, it may work, but repeating only these types of tests will leave the organization exposed to systematic issues. Evolution of testing tools: The conclusions drawn of tests are often misunderstood, security testing isn't just about what's being detected. We should think how to limit the attack paths first, rather than creating obscure, not reliable detections where log sources / monitoring tooling may not be readily available.

As an example, attacking a Crown Jewels system where detections are best suited through NDR, but not readily deployed leaves Detection Engineers with a visibility gap and consequently rely on indirect evidence in the logs with potentially many hard to action alerts. Security Testing should highlight the most critical course of action for the organization: whether the finding requires immediate infrastructure investment, preventive controls, detection engineering, or a combination of approaches. This strategic guidance transforms security testing from a compliance exercise to a risk management process which drives informed decision-making in IT and security investments.

What should be a detection

Not every security test finding should result in a detection rule. The decision to implement a detection should be guided by several criteria which balance operational effectiveness with resource constraints.
Implement detections when:

- High-fidelity signals exist - The behavior can be reliably distinguished from legitimate activity with acceptably low false positive rates
- Scalable monitoring is available - Log sources provide consistent, data across the environment
- Actionable response is possible - Security teams can effectively investigate and respond to alerts
- Business impact justifies investment - The protected assets or prevented damage warrants the detection engineering effort

Avoid detections for:

- Noisy, low-confidence indicators - Techniques that generate excessive false positives without clear tuning paths
- One-time, environment-specific findings - Issues better addressed through configuration changes or patching
- Activities with limited visibility - Behaviors where log sources are sparse, unreliable, or require significant infrastructure investment
- Post-compromise artifacts - Evidence that appears only after significant damage has already occurred

Views from Ontinue's Cyber Defense Center

Our AI supported human analysts review thousands of alerts and investigations across customer environments every month. Alongside this constant threat monitoring, at least one, often multiple red or purple team assessment is taking place each week. This provides a unique vantage point to compare how simulated attackers operate against how live adversaries behave when their goal is persistence, evasion and monetization.

Red teams deliver structured, scoped exercises intended to measure resilience and highlight areas for improvement. They frequently include phishing-resistant credential attacks, adversary-in-the-middle techniques, and Azure identity abuse. Yet they remain constrained by rules of engagement, a slow tempo marred by operational safety concerns, and limited time windows.

Real-world attackers face no such constraints. They exploit overlooked entry vectors such as SVG phishing attachments, replay stolen tokens to bypass multi-factor authentication, persist through Azure AD applications and automation accounts, and deliberately tamper with monitoring to maximize dwell time. We examine recurring patterns observed across live investigations and contrasts them with the practices commonly seen in red team exercises.

Key Findings from Ontinue Cyber Defense Center

- Over 70% of phishing attachments that bypassed secure email gateways in the last six months were non-traditional formats such as SVG or IMG (quishing), not Office macros or traditional Malware.
- Roughly one in five live incidents investigated included refresh token replay being used for persistence after a password reset, bypassing MFA.
- Nearly 40% of cases involving attempted Azure AD persistence showed adversaries layering multiple methods (application + automation job + role escalation) for redundancy.
- In intrusions where Azure persistence was established, attackers always attempted to tamper with diagnostic settings or conditional access policies.
- Median dwell time in cloud intrusions exceeded 21 days when adversaries successfully suppressed or manipulated telemetry.
- A significant proportion of incidents (over 30%) involved attackers attempting to leverage Azure control plane features (RunCommand, Data Factory, Key Vault) for exfiltration.
- Each week, at least one red or purple team exercise occurs across Ontinue customer environments, providing structured contrast to live adversary activity seen across thousands of alerts.

Phase 1: **Initial Access**

Red Teams:

Red teams commonly run controlled phishing campaigns as their entry point. Mature teams employ adversary-in-the-middle kits such as Evilginx2 or Modlishka to capture session cookies, enabling MFA bypass. Others deliver OAuth consent phishing campaigns, asking users to grant Graph API permissions to rogue apps. While email payloads are often HTML or PDF lures, some exercises now include QR code phishing ('quishing'). Safety constraints usually limit the use of executable payloads or uncommon file types like SVGs, even though these are increasingly exploited in the wild. Some customers provision test accounts to ensure success, meaning the exercise bypasses the realistic challenges of initial compromise.

Threat Actors:

Attackers increasingly abuse overlooked file types, especially SVGs. SVGs can contain embedded JavaScript or encoded redirects that trigger within the browser. Email security products often treat them as harmless images, letting them through. Once opened, the victim is redirected to an adversary-controlled Microsoft login clone. Both credentials and refresh tokens are captured. Analysis shows adversaries rapidly validate credentials against Exchange Online or SharePoint, then replay tokens via Graph API within the same day. This compresses the timeline from phishing to full tenant access to hours rather than days.

Phase 2: **Authentication & Token Abuse**

Red Teams:

Advanced red teams do not stop at username and password theft. They demonstrate MFA bypass by capturing and replaying cookies with Evilginx2 or by using TokenTactics to mint device tokens to push on for further Azure services. Some simulate refresh token replay to show risk, but sustained automation across weeks is usually avoided due to cost. Exercises often end once access is proven, rather than maintaining stealthy persistence cycles. Extended persistence via refresh tokens is rarely tested.

Threat Actors:

Real adversaries abuse refresh tokens aggressively. A single stolen token can be replayed indefinitely until revoked, minting new access tokens without fresh login events. This bypasses MFA entirely. Investigation data shows adversaries scheduling automated token refreshes every 30–90 minutes, maintaining continuous access while avoiding login telemetry. In some incidents, access continued for weeks after a user's password was reset because refresh tokens were never revoked.

Phase 3: **Persistence in Azure AD**

Red Teams:
Persistence is simulated through limited app registrations with minimal permissions. A single app secret is usually created, and escalation is avoided due to safety constraints. Red teams register Azure AD applications with limited permissions, add a single secret, and sometimes request delegated access through OAuth consent. Privilege escalation and automation job creation are typically avoided due to safety constraints, though experienced teams may use <define tools> or custom tooling to illustrate the potential for backdoored applications.

Threat Actors:
Real adversaries establish redundancy. They create Azure AD applications with broad permissions, add multiple client secrets and certificates, and assign high-value roles. Some configure automation jobs or Logic Apps disguised as business processes. By overlapping expiry dates on secrets, attackers maintain uninterrupted access even if one credential is revoked. Investigations show multiple persistence layers frequently coexisting, complicating incident response.

Phase 4: **Expansion & Control Plane Abuse**

Red Teams:
Exercises usually stop at demonstrating potential. They may retrieve one Key Vault secret, use RunCommand on a single VM, or show read access to a storage container. The purpose is to prove risk, not to operationalize exploitation. Red teams avoid scale to minimize disruption, although they document how escalation could unfold.

Threat Actors:
Adversaries attempt to fully exploit the control plane. Common patterns include:

- Using RunCommand to execute payloads across multiple VMs simultaneously.
- Extracting secrets from Key Vaults that unlock downstream systems.
- Generating Shared Access Signatures (SAS) for long-duration data exfiltration.
- Building Data Factory pipelines replicating terabytes of data to external accounts.

These techniques appear as legitimate administrative activity, requiring anomaly detection to identify.

Phase 5: **Evasion & Dwell Time**

Red Teams:
Rules of engagement generally prohibit telemetry tampering. Red teams note the risks of modifying conditional access or deleting diagnostic settings but rarely execute them. Instead, they simulate stealth by restricting activity volumes and using legitimate tooling to blend in where permitted.

Threat Actors:
Adversaries actively conceal themselves. They delete diagnostic settings forwarding logs to SIEM, modify conditional access rules to trust attacker devices, or suppress MFA challenges. This reduces visibility, extending dwell time. Cases show dwell times exceeding three weeks when monitoring suppression was effective.

Phase 6: **Endgame**

Red Teams:
For obvious reasons, engagements stop at risk demonstration: controlled exfiltration or simulated ransomware. Endgame scenarios are simulated rather than executed. Red teams exfiltrate limited sample files, capture screenshots of sensitive systems, or drop staged ransomware notes. Large-scale data theft or encryption campaigns are explicitly avoided, with potential impacts described in post-engagement reports. This often loses potency with executives as demonstration rarely equates to real work impact.

Threat Actors:
Adversaries monetize aggressively. Mailboxes and SharePoint libraries are exfiltrated, Key Vault secrets are harvested for broader compromise, and backups are copied or deleted. Access is sometimes resold to ransomware affiliates who use RunCommand to deploy encryption across Azure VMs and storage. Impact is maximized by chaining persistence, evasion, and legitimate-tool abuse.

Defender's Priorities

Closing the gap between testing and real adversary behavior requires focus on the following:

- Inspect non-traditional file types (SVG, ISO, IMG) for active content.
- Detect abnormal refresh token reuse, particularly long-lived sessions with no new logins.
- Audit Azure AD app registrations, secrets, and privileged role assignments continuously.
- Alert on large-scale RunCommand use, Data Factory replication, and spikes in Key Vault reads.
- Detect tampering with diagnostic settings and unusual conditional access modifications.
- Monitor spikes in data egress and SAS token creation.

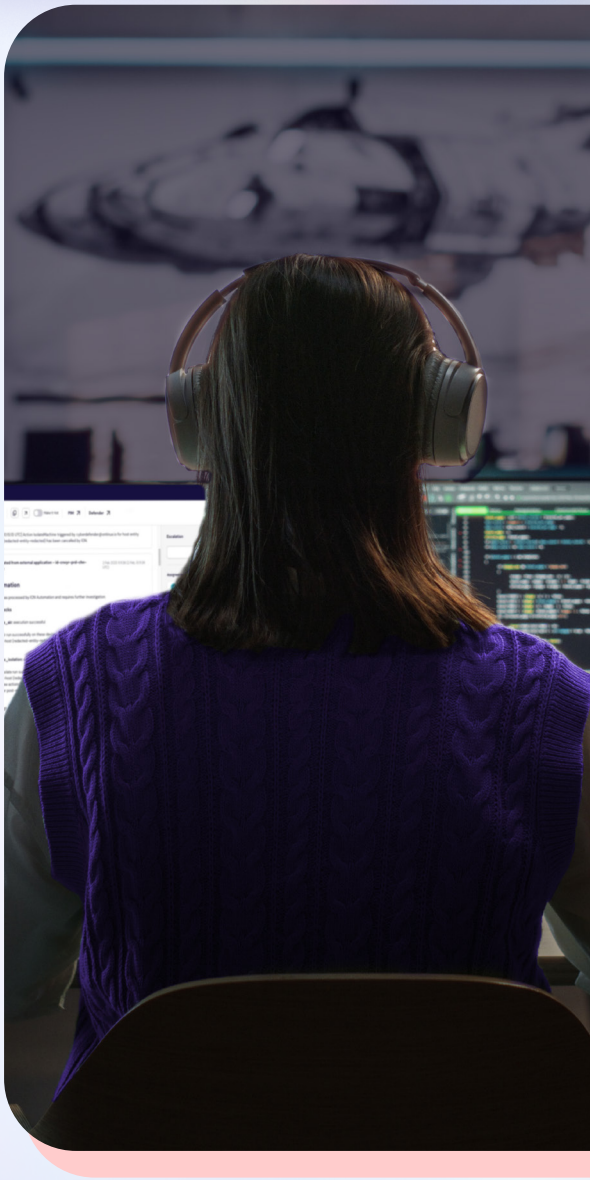

Red team testing provides an invaluable mirror, but it is only a reflection of risk under controlled light. Real adversaries operate in the shadows, with no rules of engagement, no time limits, and no regard for scope. They weaponize overlooked file types, replay tokens indefinitely, build persistence across layers of Azure, and silence the very telemetry defenders rely on.

The Ontinue Cyber Defense Center sees both perspectives daily: structured exercises that highlight where defenses should hold, and live intrusions that show where they actually fail. The insight is clear. Organizations cannot treat red team outcomes as the finish line. They must use them as a baseline, then measure against the tactics real adversaries repeatedly employ.

Red team results highlight readiness, but adversary patterns highlight reality. Persistence in Azure is no longer an exception; it is the rule. Token replay has become the modern attacker's skeleton key, bypassing MFA silently and repeatedly. The gap between simulation and reality is where dwell time grows, and it is in that gap that attackers achieve their greatest impact.

Organizations that close this red team gap by aligning purple team findings with observed adversary behavior will reduce dwell time, detect persistence earlier, and disrupt adversaries before they monetize access. Exercises validate capability. Intelligence validates resilience. Both are essential, but neither is sufficient alone.

CHAPTER 4

In the News

Geopolitics

With geopolitical tensions growing around the world in the first half of 2025, it comes as no surprise that the repercussions from these global events should also affect private sector actors, either as collateral damage or through targeted actions against private sector actors because they are seen as associated geopolitical rivals. While this phenomenon of hybrid warfare – the blurring of lines between peacetime and war with constant adversarial activity below the threshold of open warfare but meant to disrupt and unsettle an adversary - is not new, the scale, intensity, and above-all visibility of the consequences of such activity has grown increasingly in the last months. Below, Ontinue shall illustrate this with some well-known case studies.

Targeting Private Firms for Geopolitical Goals – Iran and Sepah Bank

A prime example of cyber threat actors targeting private sector entities as part of global geopolitical events can be seen with the actions of Predatory Sparrow (aka Gonjeshke Darande), a pro-Israeli hacking group, against Nobitex – Iran's largest crypto exchange – and Iran's Sepah Bank. Predatory Sparrow also claims to have been responsible for attacks against Iranian steel manufacturing facilities.

Sepah Bank, a private business, was targeted by Predatory Sparrow for its alleged funding of Iran's military. Occurring in June 2025, at a time of heightened tensions between Israel and Iran, the attack is no coincidence. Disrupting everyday life, at a time of conflict, by attacking such a key piece of national infrastructure is the perfect way for threat actors seeking geopolitical goals to keep their adversaries off-balance. The added benefit for nation states of having these actions being carried out by non-state or proxy actors, is that it allows these tensions to remain below the threshold of open warfare while still contributing towards geopolitical objectives and advantages.

By undermining a key national service, such as the banking sector, threat actors can undermine a population's trust in their governments' as a whole, as well as these governments' ability to keep their populations safe. Particularly in autocratic nations, where governments' appearance of invulnerability is vital for their survival, such a public and unavoidable mishap can have significant implications on national morale, unrelated to the public relations implications.

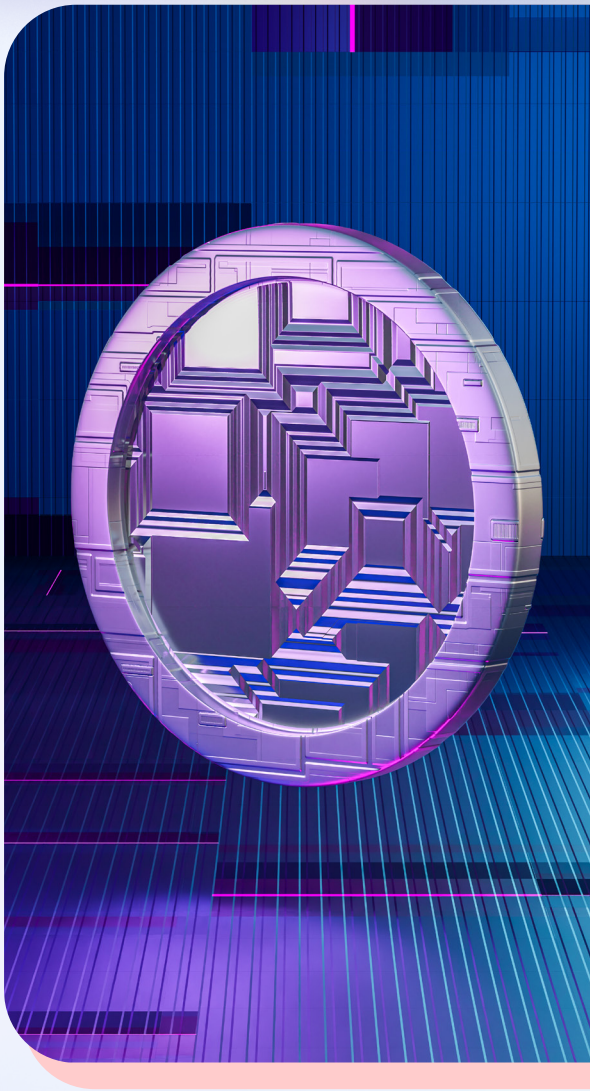

References:

- https://www.wired.com/story/predatory-sparrow-cyberattack-timeline/
- https://dailysecurityreview.com/security-spotlight/predatory-sparrow-drains-and-burns-90m-in-cyberattack-on-irans-nobitex-exchange/
- https://www.bbc.com/news/technology-62072480
- https://www.reuters.com/world/middle-east/suspected-israeli-hackers-claim-destroy-data-irans-bank-sepah-2025-06-17/

Conflict Elsewhere Leading to Increased Cyber Espionage Campaigns – Void Blizzard and the Netherlands

Ongoing military conflicts also have a way of impacting the everyday lives of average citizens, both in nations directly engaged in wars, and those that are not active participants. The war in Ukraine continues to provide numerous examples of threat actors' activities crossing into the lives non-combatants. Historically, the NotPetya ransomware in 2017, provides the most well-known accidental overspill from a hot conflict into civilian life, and similar activity continues today.

While NotPetya's global impact was unintended by its creators, in 2025, Void Blizzard or Laundry Bear are a well-known threat actor group, who deliberately target their victims, because of global geopolitical circumstances. The hacker group, which is affiliated with pro-Russian causes, were responsible for a cyber-attack against the Dutch police in 2024, and for also targeting NATO and European countries more widely. The group focuses on espionage operations, aiming to steal sensitive information about Western military equipment or arms deliveries to Ukraine. Private entities which produce advanced technologies, beyond the capabilities of Russia's native industries or unavailable for Russia to purchase due to Western sanctions, have also been known to be targets.

In the worldwide game that is strategic competition, disrupting an adversary's normal course of business can confer a strategic advantage, no matter how brief. With many Western military suppliers or critical industries primarily being run as private companies, who must focus on their shareholders interests before security necessarily, an adversary's ability to even challenge the seamless operations of these firms can have geostrategic reverberations. Since the end of the Cold War, armaments manufacturing has wound down, with fewer firms producing fewer munitions, at fewer sites. With Western and NATO munitions stockpiles already being stretched with the assistance provided to Ukraine, any halt to production schedules can have noticeable impacts on nation-states' geopolitical strategic decisions. Other considerations such as geography also come into consideration: the Netherlands' location, with the biggest port in Europe being in Rotterdam, plays a role in threat actors' decision of what or whom to target.

References:

- https://www.techzine.eu/news/security/131836/microsoft-exposes-laundry-bear-targeting-critical-infrastructure/
- https://cybercover.sg/2025/laundry-bear-unmasking-the-russian-cyber-espionage-threat-to-nato-and-european-security/
- https://www.sofx.com/russian-linked-hacker-group-laundry-bear-targeted-dutch-police-nato-networks/

Criminal Behavior by State-Sponsored Threat Actors – the Lazarus Group and North Korea

Opportunistic, criminal actions by threat actors sponsored by pariah nation-states also continues unabated. Threat actors from North Korea are the most prominent and well-known for engaging in hacking for financial gain. There is a very good reason for this: the proceeds from this criminal activity are then redirected by the North Korean state to fund their state's nuclear and missile programs.

North Korea's Lazarus Group earned their infamy with a $101-million Bangladesh Bank heist in 2016. In 2025, they were attributed with the $1.5 billion Bybit hack. The Lazarus Group have developed a reputation for successfully targeting the cryptocurrency industry. Bybit, a Dubai-based cryptocurrency exchange, fell victim to the Lazarus Group in February 2025, resulting in what is believed to be the Lazarus Group's biggest heist against a single firm to date. By leveraging the SafeWallet interface used by the exchange's executives, the hacker group executed fraudulent transactions, before distributing these across multiple wallets.

In 2023, a UN report estimated that North Korea's cyber-attacks had earned the regime approximately $6 billion between 2017-2023, with as much as 40-50% of this sum being used to directly fund the country's nuclear weapons program. Given the group's propensity for opportunistic targeting, no private firm is safe from attack. The added risk for the victim of potentially engaging and paying a ransom to a UN-sanctioned entity, only increases the perils related to compliance for the affected firm. This can strengthen the attacker's hand in negotiations and

can actually lead to firms being less transparent in terms of reporting when they fall victim to such attacks, especially when national regulations do not exist to enforce the disclosure of such attacks. The criminal nature of this activity and North Korea's continued status as a pariah state, mean that the risks associated with this threat are unlikely to reduce any time soon.

References:

- https://cointelegraph.com/learn/articles/lazarus-group-hackers-behind-billion-dollar-heists
- https://www.radware.com/cyberpedia/ddos-attacks/the-lazarus-group-apt38-north-korean-threat-actor/
- https://www.dw.com/en/how-crypto-heists-help-north-korea-fund-its-nuclear-program/a-68669802

Predictions

In the persistent game of cat-and-mouse between threat actors and cybersecurity teams, the geopolitical threats mentioned above will not dissipate in the short-term. The phenomenon may be carried out by different threat actors from time to time, but the underlying reasons for geopolitics playing a role in cyberattacks will remain constant.

Despite Iran's military leadership, organization, and offensive cyber abilities having been severely affected in the aftermath of the bombardment of the country's nuclear installations, we can venture to predict an increase in primarily nuisance attacks (e.g.: DDoS, website defacement, etc.) against Iran's perceived enemies in the shorter-term. In the medium-term, it would not be surprising for Iran's enemies to see more serious attempts to cause concrete damage against their critical infrastructure. There are precedents that Iran may seek to replicate, such as a hacker's attempt to manipulate the systems controlling the water treatment plant at Oldsmar, in the US state of Florida, in 2021. Iran also has their own alleged precedent from 2020, when Iran was accused of having launched a failed cyber-attack against Israel's water infrastructure systems to interfere with chlorine levels.

A second prediction is that hackers are increasingly going to leverage Open-Source Intelligence to target employees working at the firms they seek to target. Despite the prevalence of cyber security awareness campaigns for employees, very few provide guidance or raise attention to the risks associated with what employees post online. This is an incredibly difficult aspect for firms' security teams to monitor, and privacy regulations hinder any concrete action that can be taken to prevent an employee posting material that could be useful for threat actors on their personal profiles. One need only to read about the bodyguards to Swedish Prime Minister, Ulf Kristersson, uploading the details of their running and cycling routes to the fitness app, Strava, to see that this is an issue that can concern the most benign activity, at every level of society, and yet confer significant advantages and insights to geopolitical rivals.

Third Party Risk

An organization may invest heavily in strengthening its own security posture to prevent data breaches; however, those efforts can be quickly undermined if third-party partners with weaker security controls are granted access to its systems or entrusted with sensitive data. Data breaches attributed to third-party vendors or partners doubled globally from 2024 to 2025. In Verizon's Data Breach Investigations Report (DBIR), third-party involvement was identified in about 30% of all breaches analyzed in 2025, up from roughly 15% in 2024 .

Since the start of 2025, there has been a noticeable rise in breaches where third-party vendors have served as the initial access vector or failed to adequately safeguard the data they handle, making them the root cause of compromise.

In April 2025 in the UK a retail giant, Marks & Spencer, became victim to a catastrophic ransomware attack pulled off by Scattered Spider which resulted in a £300 million loss of revenue and £1 billion wiped off the corporation's market value. To gain initial access, Scattered Spider impersonated M&S employees and social-engineered Tata Consultancy Services helpdesk staff into resetting passwords for high-privileged accounts. At least two TCS employee accounts with legitimate M&S system access were then used during the intrusion, allowing threat actors to bypass multifactor authentication and internal controls

In May 2025, sportswear giant Adidas disclosed a cyberattack that exposed customer data through an external customer service provider. The breach exposed customer contact information that had been provided to Adidas's help desk in the past. The stolen data consists primarily of personal contact details. This includes customers' full names, email addresses, phone numbers, physical mailing address, and dates of birth.

These brands represent just a few among a growing list of organizations that have suffered data breaches where a third party was the root cause, either by storing sensitive data on their behalf or by serving as an unintentional backdoor and initial access point into their IT estates.

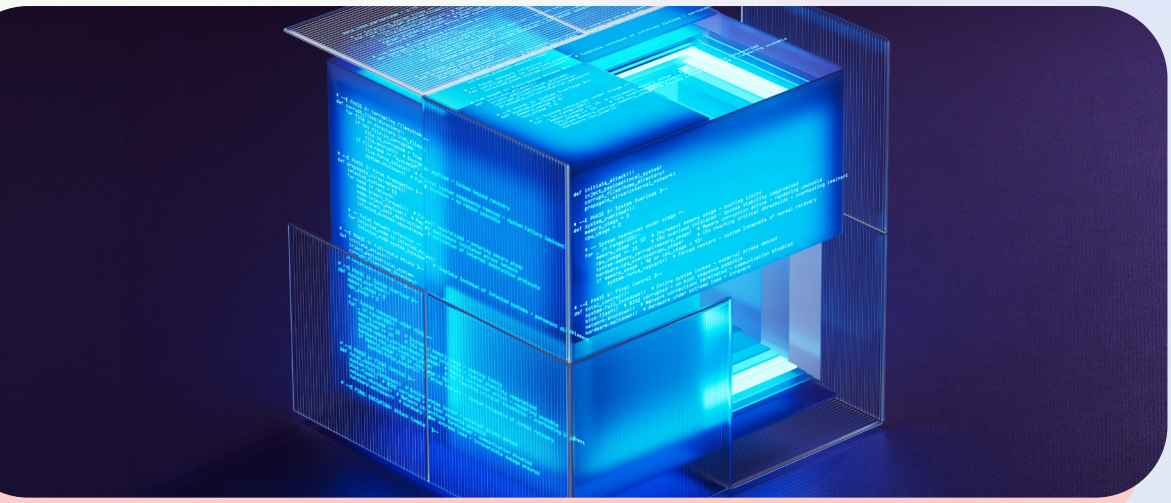

Lessons Learned

Organizations seeking to reduce costs by outsourcing help desk functions or other repetitive tasks overseas must carefully weigh the security risks associated with granting third parties access to their networks and data. It is critical to factor in the potential impact of data breaches when evaluating such partnerships. To manage this risk effectively, organizations should establish robust process frameworks that identify all third-party relationships, clearly define the systems and data each vendor can access, and implement thorough vetting procedures to ensure each partner maintains adequate security controls.

Deepfakes

Deepfakes present an ever-growing problem in the digital landscape. So far about 75% of all Deepfake incidents were targeting either public figures – both non-consensual explicit content or other means leading to financial extortion as well as misinformation /influence operations – or scamming / extorting public citizens.

Almost one fifth of the total incidents targeted organizations, mostly financially motivated.

Impersonation for credential theft via deepfakes is relatively novel, but it's a trend worth watching for, since the one of the first notable incidents: Hong Kong CFO deepfake video scam resulting in a $25M transfer in 2023.

Notable Incident types:

- Romance scams (also known as: "pigbutchering")
- Celebrity Investment Scams (cryptocurrency / gift card etc.)
- Fake Events Scams

Ontinue estimates that in 2025 alone we will see a total of about 500 documented deepfake incidents.

Source: Resemble AI Deepfake Incident Report.

Emerging Social Engineering Trends

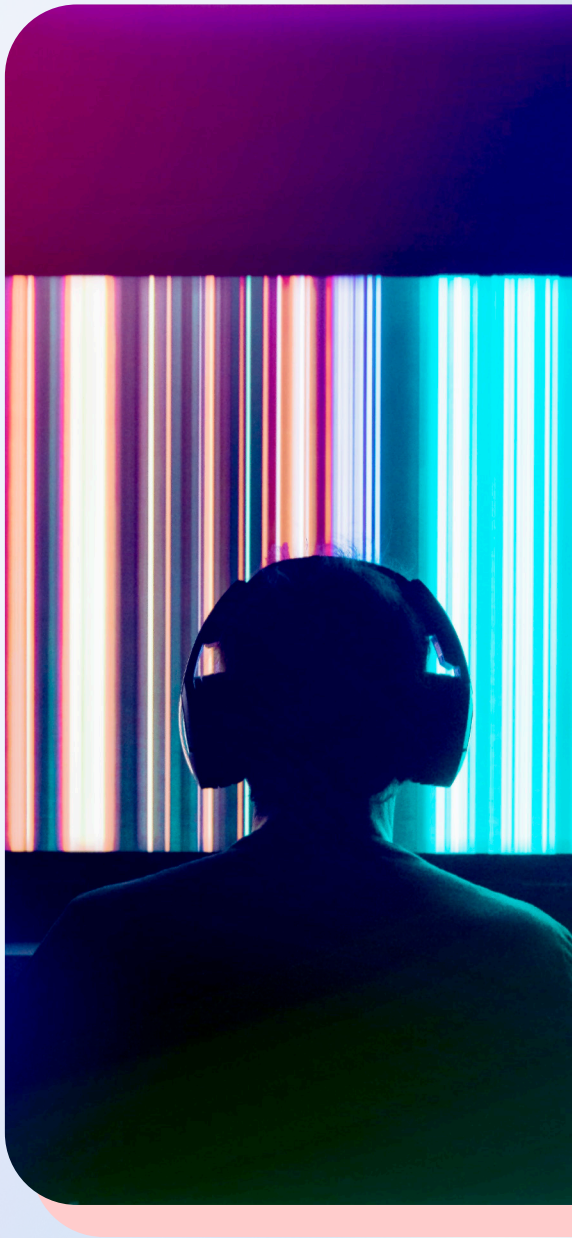

Navigating information in our digital lives have never been more difficult – considering both volume and variety. Statista.com estimates over 400 million Terabytes of data created every single day in 2025. Both professional and social networks suffer from information overload and misinformation.

While "digital natives" may excel at handling volume, distinguishing reality from half-truths and complete fabrications have become increasingly technologically and mentally taxing – practicing critical thinking constantly does not come organically.

Evolutionarily we have not adapted to operate in a trust-eroded misinformation era, where we need to scrutinize our senses steadily – losing the link between our perceptions and reality also creates lurking anxiety.

In the past decades, most educators celebrated the internet for providing access to all human wisdom, "more information", exposing us to diverse opinions – only to recently turn into algorithm driven self-perpetuation and opinion echo chambers.

Enter the Generated Information Age

It appears that parts of the "dead internet theory" are becoming reality – namely that organic human activity is being replaced by bots - describing a negative sentiment of the generative AI age.

In other words, most new content on the internet, trends, content is and will be generated rather than human captured or curated - disintegrating and making our approximation of objective reality vulnerable and undermining genuine digital human to human interactions (especially with strangers) – we will need to arm ourselves with new safety mechanisms.

Still valuable communities are oftentimes becoming paywalled, invite-only, and part of the deep web instead of the surface web - where algorithms, hooks and thumbnails define popularity and human attention is steadily becoming the primary internet currency, while on the individual level we have a dwindling amount to offer.

Consequently, everyone starts to understand as information is becoming less trustworthy - and given the difficulty of verification - there is an increasing importance of the emergence of post-Captcha verification mechanisms to defeat bots, validate personal digital identities, all of which is potentially leading to privacy and anonymity undermining solutions being rushed by legislators.

Neither civilians nor corporate employees are prepared – regardless the amount of Social Engineering trainings. How often have we heard, after all about Security Professionals admitting also falling for convincing Phishing emails? Consequently, how many times do we think we are being misled on a daily basis, given the allotted short attention. How often are we stopping anyway for critical evaluation? Are we in the habit of rushing to consume information while fearing of missing out?

Naturally, in the Information Age there was no shortage of mis- / disinformation either – a trend that continues to grow steadily starting from past centuries. Information however is becoming unverifiable in all formats: text, web, image, audio and video formats all suffer from validation issues being rushed.

Evolved from simple phishing emails and websites, we see an increasingly rich media: images (SVG Smuggling), audio (voice cloning), videos (deepfakes) and conference calls (face swapping/cloning) all emerge as a valid attack vector.

One example of Fake Identities combined with persistent access is to North Korean Remote IT workers using made-up profiles and connecting from Laptop Farms.

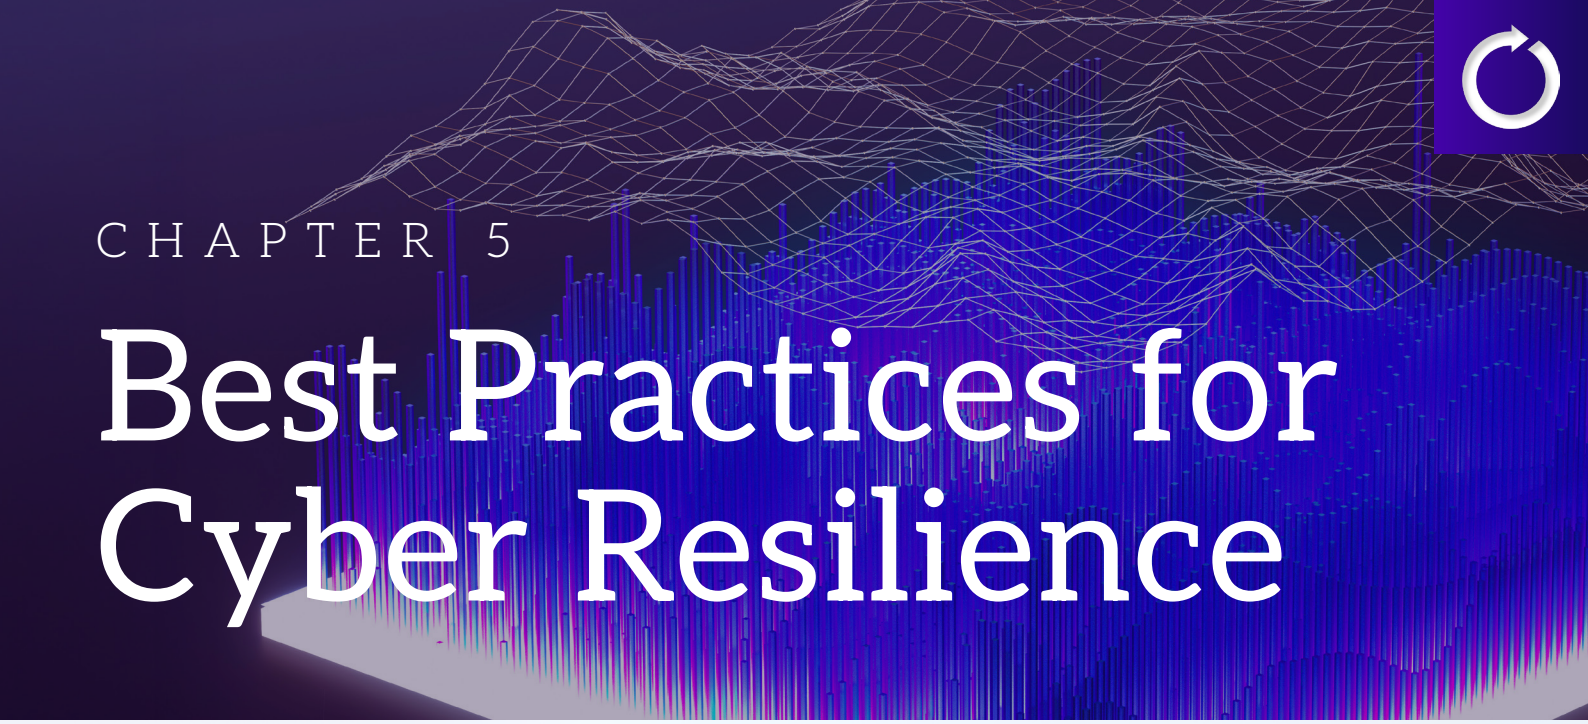

CHAPTER 5

Best Practices for Cyber Resilience

The first half of 2025 reinforced that while attackers are innovating, many successful breaches exploit overlooked basics. Organizations can strengthen resilience by focusing on the following priorities:

1. **Prepare for Ransomware:** Maintain offline, tested backups and patch high-risk vulnerabilities quickly.
2. **Protect Identities:** Deploy phishing-resistant MFA and monitor for token theft or session hijacking.
3. **Control Endpoints:** Enforce least privilege, restrict USB use, and enable EDR/XDR in blocking mode.
4. **Manage Third-Party Risk:** Require vendors to meet baseline controls and include them in incident response plans.
5. **Counter Social Engineering:** Train staff to verify unusual requests, watch for deepfakes, and reinforce phishing awareness.
6. **Prioritize Intelligence-led Defense:** Use threat intelligence to guide testing, red-teaming, and detection tuning.

By focusing on these fundamentals, organizations can reduce the likelihood of compromise while improving their ability to contain incidents when they occur.

The Role of MXDR

Even the strongest security strategies require consistent execution and rapid response. An MXDR partner augments internal teams by:

- Monitoring threats 24/7 across endpoints, identities, and cloud.
- Operationalizing threat intelligence to detect real-world adversary behaviors.
- Providing expert human response alongside AI-driven automation.

Working with a strong managed security partner will help organizations stay ahead of fast-moving threats so CISOs and their teams can focus on bigger business priorities.

Contributors

Ontinue's Advanced Threat Operations (ATO) team leverages proactive threat identification, analysis, and mitigation to empower our customers with the resilience needed to tackle the constantly evolving threat landscape.

Special thanks to:

Will Bailey - Senior Cyber Defender

Kaan Vartuk - Cyber Advisor

David Reich - Senior Detection Engineer

Adam Bennett - Principal Cyber Defender

Usha Sree Yannapu - Senior Cyber Defender

Ontinue

Ontinue offers nonstop SecOps through an AI-powered managed extended detection and response (MXDR) service. Ontinue ION MXDR combines powerful proprietary AI with the industry's first collaboration with Microsoft Teams to continuously build a deep understanding of our customers' environments, informing how we prevent, detect, and respond to threats.

Continuous protection. AI-powered Nonstop SecOps. That's Ontinue.

CONTACT US

LEARN MORE