JULY 2026

FROM THE FRONT LINE

Q2 THREAT UPDATE

CYFOR SECURE
CYBER SECURITY

What We're Seeing

Daniel Whelan | Front Line Perspective

Daniel Whelan is a Cyber Incident Response Recovery Specialist at CYFOR Secure, working as part of our cyber incident response team to help organisations recover from ransomware and other cyber attacks. Supporting businesses through every stage of the recovery process, he brings practical, frontline insight into the challenges organisations face after an incident.

Three months on from our last update, the picture has not improved. If anything, the trends we flagged in April have hardened into operating norms.

The ransomware ecosystem is no longer fragmenting. It is consolidating around a smaller number of more capable groups, and those groups are openly trading affiliates, tooling and infrastructure with each other. The labels on a leak site post tell us less about who actually did the work than they did a year ago.

The economics have shifted too. Fewer victims are paying, so the groups that remain are extracting more from each engagement and increasingly skipping encryption altogether. We are now investigating incidents where there is no encryptor at all, just a quiet exfiltration, a leak site listing, and a demand. That removes the backup-as-saviour argument from the table.

Phishing has come back as the dominant initial access route. It is not the phishing of two years ago. It is OAuth device code abuse, session token theft, and adversary-in-the-middle kits sold for a few hundred dollars a month. The MFA you rolled out in 2022 will not stop most of what we are see-

-ing on the front line.

Two strong arrests of suspected Scattered Spider members in April have done little to slow the playbook those operators pioneered. Other crews have picked it up. Help desk impersonation, MFA reset social engineering, and voice cloning are now common across multiple unrelated groups.

This update covers what our team has seen and worked across April, May and the early weeks of June 2026, with a look ahead at Q3. The takeaway is the same one we wrote last quarter, only sharper: prevention alone is not enough. Assume exposure, plan for exfiltration, harden your help desk, and patch the things on the front of every advisory list.

The Quarter in numbers

A handful of figures worth holding in your head, and a few comparisons that show how fast Q1 has moved into Q2.

Leak site volumes. Q1 2026 produced around **2,122** known ransomware victims posted to leak sites. April alone produced around **770**, putting Q2 on pace to match or exceed Q1. The summer slowdown that has been a reliable seasonal pattern for years is showing no sign of arriving.

Cartel share. The top ten ransomware groups now account for roughly **71 percent** of all leak site victims. At the end of last year, that share was significantly lower. The fragmentation trend has reversed inside two quarters.

Group movements. Qilin held the top spot for the fourth consecutive month. LockBit, rewritten and re-released as 5.0, jumped from **79 victims** in Q4 2025 to **163** in Q1 2026, an increase of more than 100 percent quarter on quarter. The Gentlemen, almost unheard of in late 2025 with **40** Q4 victims, climbed to **166** in Q1 and added more than **80** again in April alone. DragonForce posted **101** in Q1 and **63** more in April, on track for around **190** across Q2. Akira moved the other way, down **22 percent** QoQ at **176** in Q1 and only **48** in April.

Payment behaviour. The publicly tracked ransom payment rate fell from around **63 percent** in 2024 to around **28 percent** across 2025, with Q4 2025 figures sitting closer to **20 percent**. Q1 and Q2 2026 numbers are not yet published but the trajectory is unambiguous.

UK exposure. Around **204** cyber attacks were assessed as nationally significant in the year to August 2025, more than **double** the previous twelve month period and equating to roughly **four** every week. Around **98 percent** of UK higher education institutions identified a breach or attack in the last twelve months, up from **91 percent** the year before.

98%

of UK higher education institutions identified an attack or breach in the last 12 months

Up from 91% the year before

Initial access. In the majority of ransomware engagements our team has worked across Q2, the victim's credentials were already in infostealer logs or initial access broker listingss before the encryptor or the leak site listing ever appeared.

Ransomware has consolidated, not calmed down

The headline statistic for Q2 is that the top ten ransomware groups now account for roughly 71 percent of all victims on leak sites. Twelve months ago that share was much lower as smaller crews splintered off from Conti derivatives and tried their hand at running their own brands. That experiment has largely failed. Smaller groups have not generated meaningful affiliate interest, and the larger groups have absorbed the talent.

The Gentlemen, which barely existed at the end of last year, has become the fastest scaling group on record. We worked one Gentlemen engagement in Q2 against a UK software consultancy whose access was then re-used to pivot into one of its clients overseas. The group then publicly labelled the UK firm an access broker as part of the pressure campaign. It is a tactic we will see again. The reason that matters is that The Gentlemen are not just giving affiliates an encryptor. We have also evidenced the operators packaging the supporting tradecraft too: EDR killing tools, vulnerable driver abuse, multi-OS lockers and pivot infrastructure designed to get affiliates to the encryption or extortion stage quicker with less chance of early detection.

LockBit 5.0 has come back faster than anyone hoped. After its near collapse following the law enforcement action of last year, the rewrite is faster, more modular, and runs across Windows, Linux and ESXi.

Underneath the brands, the cartel structure is now openly acknowledged. LockBit, Qilin and DragonForce announced their alliance last autumn and have been operating with shared infrastructure and affiliate pools through Q2. A second consortium publicly branded as

LAPSUS$ Hunters, or under its Trinity of Chaos badge, brings together ShinyHunters, LAPSUS$ and operators historically known as Scattered Spider. They run what amounts to an extortion as a service operation with paying affiliates, and they have layered on aggressive personal harassment tactics including swatting of executives, contacting journalists and notifying regulators on behalf of their victims.

If you refuse to pay one group, your data may still appear under a different banner. That is not theoretical anymore.

ShinyHunters owns the Quarter

If one group has defined Q2 2026, it is ShinyHunters. They moved from being a recognisable leak site brand to being the most active extortion operation in the world, with more than 40 victims listed on their site since the start of the year.

On 17 April, they listed Medtronic, claiming over nine million records. Medtronic confirmed the unauthorised access in an SEC filing on 24 April but disputed the volume; the listing later disappeared from the leak site without a public dump, a pattern we associate with private negotiation. On 18 April, they listed Carnival Corporation after social engineering of a single employee account on 14 April; Carnival began customer notifications on 27 May to just under six million affected individuals. On 20 April, they exploited stolen authentication tokens from ADT and dumped 11 gigabytes the following week after ADT refused to pay. On 11 April, they claimed Rockstar Games via the Anodot and Snowflake supply chain, threatening over 78 million analytics records. Charter Communications, Pitney Bowes, McGraw-Hill, Mytheresa, Zara, 7-Eleven and Udemy all appeared on the same site through April and May. On 1 May, they hit Instructure (covered separately under Education).

The pattern across every one of these engagements is the same. Initial access is achieved through vishing or token theft against a single user, often with downstream connections to Salesforce, Snowflake or Microsoft 365. There is no traditional encryptor. There is no ransomware binary detonated against production. There is simply quiet exfiltration, a listing on the leak site, and a deadline. ShinyHunters wants money and does not need to disrupt operations to get it.

That model is now spreading. Pure extortion incidents (data theft only, no encryption) made up a growing share of the work we did this quarter. A clean and tested backup will help you recover an encrypted environment. It will not stop your customer data from being sold on a leak site. The conversation we are having with clients has changed accordingly.

Education: a five alarm Quarter

Education has had its worst quarter on record, and the timing was not a coincidence. UK and global exam season concentrates both the value of the data sitting in these systems and the operational leverage available to an attacker who can take them offline.

The standout incident was the breach of Instructure, the operator of the Canvas learning management platform used by thousands of UK, US, Canadian, Australian and other institutions. Unauthorised access began on 25 April and was detected on 29 April. Instructure was disclosed publicly on 1 May. ShinyHunters claimed responsibility on 3 May and demanded a ransom by 12 May. On 7 May, before that deadline expired, Canvas login pages across roughly 330 institutions were defaced with a ransom note. The platform was taken offline. Instructure announced an agreement with the actor on 11 May and stated that the data had been destroyed.

ShinyHunters claim 3.65 terabytes of exfiltrated data covering around 275 million users across 8,809 institutions. That figure should be treated as the actor's claim rather than as a verified count, but the institutional footprint is real and includes UK universities. Confirmed exposed data includes names, email addresses, student IDs and private messages between students and staff. Final exam timetables at several institutions were disrupted.

"...if your VPN does not enforce MFA, you may already be listed..."

Closer to home, AQA, OCR and Pearson Edexcel have all been hit by computer misuse incidents this quarter, with exam papers reportedly extracted and offered for sale online. AQA secured a High Court injunction against further distribution. Cambridgeshire Constabulary, Surrey Police, the NCA and the Department for Education are involved. The reported initial vector in at least one case was a compromised school email account used to make plausible-looking requests to the boards.

Behind these incidents sits a broader set of numbers. The UK government's Cyber Security Breaches Survey, published at the end of April, found that around 98 percent of higher education institutions identified a breach or attack in the last twelve months. Twenty-nine percent reported attacks at least weekly. Globally, attacks on the higher and further education sector rose by around 63 percent year on year, with FunkSec (an AI-assisted ransomware operation that emerged in late 2024) accounting for around 23 percent of observed ransomware activity in the sector. Cl0p remains the most financially damaging, with reported

demands against education targets well above ten million dollars.

For UK schools, colleges and universities heading into the next exam cycle: pressure test your SaaS suppliers on patching and SSO posture before the next window opens, pre-arrange an incident response retainer, and plan for a learning platform outage in the same way you plan for power loss.

When ransomware reboots the defenders out of the way

- from our own analysis

From static and dynamic malware analysis we performed in Q2, two TTPs keep co-occurring in our ransomware casework this quarter: Bring Your Own Vulnerable Driver (BYOVD) and forced reboot into Safe Mode with Networking. Same goal - silence the EDR before encryption - attacked from opposite sides of the boot.

Every modern business has an "antivirus on steroids" agent (EDR) watching every program that runs. Both of these techniques are different ways of switching that watcher off long enough to encrypt every file in the company. One does it on a live machine; the other does it by restarting the machine into a stripped-down version of Windows where the watcher never wakes up.

The Gentlemen activity is the same trend at RaaS scale. Their affiliate tooling includes GentleKiller and other EDR-killers such as HexKiller, ThrottleBlood and HavocKiller, using BYOVD tradecraft to load vulnerable or abused kernel drivers and terminate security processes before payload deployment. Binaries such as Sent.exe, Kasps.exe, Valorant2.exe and Sophos.exe, and driver artefacts such as ThrottleBlood.sys, vgk.sys, G11.sys and havoc.sys, should therefore be treated as more than standalone IOCs: they are indicators that the attacker is in the pre-encryption phase and actively trying to blind the endpoint stack.

A recent Qilin (Agenda) sample we triaged makes the Safe Mode trick almost mundane: a freshly Rust-compiled 4.69 MB Windows binary, unsigned, with zero VirusTotal / Malware Bazaar / NSRL / CIRCL hits at time of analysis. Inside its strings sit a dedicated [INFO|SAFE] log channel and the literal command bcdedit /set {current} safeboot network, invoked under runas via BCDEdit.exe. Pre-stage the next boot into Safe Mode with Networking, drop a RunOnce to re-launch the encryptor at login, reboot, and finish encryption in a Windows session where most third-party EDR drivers - registered SERVICE_AUTO_START against the normal boot but not listed under HKLM\System\CurrentControlSet\Control\SafeBoot\Network - never load. Network stays up, so the encryptor can still hit SMB shares and call vCenter; the watchdog is asleep.

What that means in practice is that the malware tells Windows “next time you start, start in

Safe Mode but keep the network on", schedules itself to relaunch automatically after that restart, then forces the reboot. Safe Mode is the cut-down Windows you've probably used to fix a broken PC - Microsoft only loads a short list of approved drivers in that mode, and most third-party security tools are not on the list. The criminal has now arranged for a Windows session that has internet access but no antivirus, and walks straight in to encrypt your files and reach across to your servers.

The same binary also embeds a full Windows CryptoAPI / BCrypt encryption pipeline, a custom PsExec service for SMB lateral movement (the developer's PDB paths leaked into the binary as …\psexec\…\psexec.pdb and psexesvc.pdb), and a PowerCLI module that connects to vCenter, disables HA/DRS on every cluster, changes the ESXi root password, enables SSH, and uploads the payload to each hypervisor - "Uploading and executing payload on all ESXi hosts in current vCenter", in the operator's own log strings. Ransom note signed -- Qilin. The malware names itself; it also documents itself.

The same file carries everything it needs to do the rest of the job. It can scramble your files using Windows' built-in cryptography functions. It carries its own copy of a Microsoft sysadmin tool ("PsExec") to jump from one Windows machine to the next across the office network. And it carries a PowerShell script written specifically to log into VMware vCenter - the brain that runs all your virtual servers - change the admin password on every virtual-server host, switch on remote login, and push the ransomware out to each one. That last step is the worst-case scenario: every virtual machine in your data centre encrypted at the same time, from a single compromised admin laptop. The author was sloppy enough to leave a build-system breadcrumb (…\psexec\…\psexec.pdb) inside the file and to sign their own ransom note -- Qilin, which is how we knew exactly who built it.

BYOVD is the live-system half of the same kill chain: drop a legitimately-signed-but-vulnerable driver (historical favourites: dbutil_2_3.sys, gdrv.sys, aswArPot.sys, kprocesshacker.sys, Truesight.sys), load it with sc create or NtLoadDriver, and use its IOCTL surface to terminate the EDR service from kernel context where its self-protection cannot reach. We didn't observe BYOVD in this specific build - execution was gated behind an affiliate password and the destructive path was never reached - but it's standard Qilin loader-stage tradecraft and it pairs naturally with the Safe Mode reboot: one kills the agent now, the other guarantees it stays dead through the encryption pass.

BYOVD ("Bring Your Own Vulnerable Driver") exploits a quirk of how Windows trusts software. Drivers - the low-level code that runs underneath Windows itself and has more authority than your antivirus - must be digitally signed by their vendor. Attackers don't crack the signing; they look for old, signed-and-still-valid drivers with known bugs in them, carry one with them into the victim's network, install it like a legitimate driver, and then use the driver's own bug to kill the antivirus from a privileged position the antivirus literally cannot defend against. We didn't catch this build doing it because the file requires a secret password to run its destructive code path and we didn't have that password - but every public report on this gang says they do, and the two techniques work hand-in-glove: BYOVD silences the antivirus on the live system; the Safe

Mode reboot makes sure it stays silenced for the actual encryption.

Mitigations, cheapest first. Turn on Microsoft's vulnerable-driver blocklist (HVCI / WDAC) - Microsoft maintains the list, you just enrol, and it kills most commodity BYOVD overnight.
Turn on Microsoft's free, auto-updating list of known-bad drivers - that single switch (HVCI / WDAC) blocks most BYOVD attempts. Ask your security team to check that your antivirus is configured to run in Safe Mode as well as in normal Windows; many aren't, by default, and that is the gap the malware is aiming for. Treat the commands that ransomware uses to tamper with the boot (bcdedit) and to delete backups (wbadmin, vssadmin) as alarm-worthy events on production servers - there is essentially no legitimate reason for them to fire. Protect your VMware vCenter the way you protect your domain admins: multi-factor authentication, separate admin accounts, and alerts on the specific actions the malware needs to take (mass password resets across ESXi hosts, turning on remote login across all hosts, forcing cluster shutdowns). And do not let unsigned, never-before-seen 4-megabyte executables run out of %Temp% - application allow-listing on user-writable folders shuts the front door.

Technical review:
Audit HKLM\System\CurrentControlSet\Control\SafeBoot\{Minimal,Network} and confirm your EDR's service entry is present in both; if it isn't, that one registry entry is the difference between running and not running in the encryption window. Block bcdedit.exe, wbadmin.exe and vssadmin.exe for non-admins and SIEM-alert on any bcdedit /set {…} safeboot … invocation - there is no legitimate reason for that command on a production server. Treat vCenter as Tier-0: MFA on SSO, alerts on Stop-Cluster -Force, on bulk ESXi password changes, and on bulk SSH-enable events. Application allow-listing in %Temp%, %AppData% and the per-run %Temp%\7zO*\ self-extract folders closes the executable drop path. Secure Boot + BitLocker (TPM+PIN) + a SIEM rule on Kernel-Boot Event ID 27 and on any change to HKLM\System\CurrentControlSet\Control\SafeBoot catches the boot tampering on the way in - not after the files are already encrypted.

Phishing is back, in a form your filters will not catch

Three months ago, we wrote that vishing was overtaking email phishing as the dominant social engineering vector. That is still true, but email phishing has not gone away. It has come back in a form that the traditional defences struggle with: token theft and OAuth abuse.

On 4 May, Microsoft's Defender Research team published details of a campaign run between 06:51 UTC on 14 April and 03:54 UTC on 16 April. In that single 45-hour window, over 35,000 users across more than 13,000 organisations in 26 countries were compromised through a coordinated adversary in the middle operation. The lure was a fake employee code of conduct

or HR disciplinary PDF. The flow captured both credentials and a live session token. Standard MFA was bypassed because the attacker simply replayed the token from their own infrastructure.

On 21 May, the FBI issued a public service announcement warning about Kali365: a phishing as a service kit, sold through Telegram for 250 dollars a month per tenant or 2,000 dollars for a year. It abuses the Microsoft OAuth 2.0 device authorisation flow to capture access tokens without ever needing to intercept a password or an MFA code. It ships with an AI-generated lure library and a real-time dashboard for tracking compromised users. The barrier to entry for a credible AiTM campaign is now lower than the cost of a software seat.

Sitting alongside Kali365 is EvilTokens, another device code abuse platform that became publicly tracked in February and was being widely deployed by the end of March. Refresh tokens captured by these kits remain valid for up to 90 days and survive password resets because the underlying Primary Refresh Token is what is being stolen.

The long-established Tycoon 2FA kit was taken down in a Europol-coordinated operation in March, with more than 300 domains seized. That is a real success but it has not reduced kit-driven phishing volume; Mamba 2FA has effectively doubled in size and is now the dominant offering, and total kit-driven phishing is higher today than before the takedown.

The defensive response that actually works sits in Entra Conditional Access. Token Protection cryptographically binds a token to the device it was issued to and is the single most useful new control we are recommending across Q2 engagements. Disable OAuth device code flow unless you have a documented business need. Block legacy authentication. Move privileged accounts to phishing-resistant MFA (FIDO2 or Windows Hello for Business). Train users on what a legitimate device code prompt looks like, because the existing 'check for the padlock' training does not cover this attack pattern.

The Scattered Spider arrests have not stopped the playbook

Two significant arrests have landed in Q2. On 10 April, Peter Stokes, a 19-year-old US and Estonian dual citizen, was arrested at Helsinki Airport while attempting to board a flight to Japan. Charges had been filed under seal in the Northern District of Illinois in December 2025 and were unsealed after the arrest. He remains in Finnish custody pending US extradition. On 17 April, Tyler Robert Buchanan of Dundee pleaded guilty in a US federal court to conspiracy to commit wire fraud and aggravated identity theft. He was extradited from Spain in April 2025 and faces sentencing in August.

Our own perspective on this group has sharpened over Q2 through CYFOR Secure's active engagement as expert witnesses on two separate matters involving operators connected to the broader Scattered Spider cluster. Without going into detail on either matter, the work has given our investigators a first-hand view of how the group structures its tradecraft, recruits and rotates affiliates, and selects its targets. The picture that public reporting has begun to describe (loose, English-speaking, highly mobile, and increasingly multi-branded) is consistent with what we are observing on the inside of these investigations. That direct experience now shapes the way we triage suspected Scattered Spider activity for our wider client base.

What is striking is how little the arrests have slowed the underlying playbook. There are no new intrusions publicly attributed to the historical core of the group since the arrests landed, but the techniques those operators pioneered (help desk impersonation, MFA reset social engineering, vishing of identity providers, abuse of SSO chains to reach Salesforce, Snowflake and Microsoft 365) are now standard practice across multiple unrelated crews. ShinyHunters is the most visible adopter. The same TTPs are visible in their Q2 campaigns against ADT, Charter, Pitney Bowes and Carnival.

The UK retail wave of M&S, Co-op and Harrods from April and May 2025 remains the case study for what this playbook can do to a mid-market organisation. M&S has now confirmed an operating profit impact in the region of 300 million pounds for its 2025 to 2026 financial year before mitigations, with around 100 million pounds recovered through insurance. The Cyber Monitoring Centre formally categorised the M&S and Co-op events as a single Category 2 systemic incident with a total impact between 270 and 440 million pounds.

£300m

Estimated profit impact of the M&S cyber attack before mitigations

£100 million recovered through insurance

The lesson for the rest of us is direct: the help desk is now the front door. If your service desk will reset a password or re-enrol an MFA factor on the strength of a confident phone call and a few personal details that the attacker has already scraped from LinkedIn, you are exposed. Implement out-of-band callback verification for all privileged account changes, require a manager approval step for MFA resets, and brief help desk staff that voice cloning is now a default capability rather than a novelty.

Dark Web and External Exposure Monitoring

A significant focus of our work across the last two quarters has been the rollout of a continuous dark web and external exposure monitoring service for retained clients, with a client-facing dashboard that surfaces issues in close to real time.

The capability ingests data from cyber crime forums, ransomware leak sites, paste sites, Telegram channels, criminal marketplaces and stealer log dumps, and continually correlates that data against the client's domains, employee email addresses, brand assets, executive personas, source code repositories and exposed infrastructure footprint.

Across Q2 the service has produced three categories of findings for clients.

First, credential exposure. Where employee credentials appear in stealer log dumps or combination lists, the client receives an alert with the source device, the data dumped, and a recommended response window. Several engagements this quarter have begun with a stealer log alert on the client's dashboard rather than with an incident; credentials were rotated and the source device remediated before they were ever used.

Second, infrastructure exposure. Where an initial access broker lists what appears to be access to the client's environment (typically advertised by access type, geography and sector on Russian language forums), the client is notified and our investigators fingerprint the listing to confirm whether it matches the client estate. Multiple Q2 engagements have been pre-empted this way before access was sold on to a ransomware operator.

Third, brand and executive impersonation. Where attacker infrastructure is configured to spoof the client (phishing pages, fake executive social profiles, look-alike domains, deepfake-ready collateral), the alert is paired with takedown action.

The strategic value is visibility over the part of the attack chain that sits outside the perimeter, where most precursor activity to a serious incident now lives. The majority of the ransomware engagements we worked in Q2 had observable precursor activity in the underground that would have been visible to a continuous monitoring capability. Knowing about it days or weeks before the incident lands changes what an incident response engagement looks like, and in several Q2 cases has prevented the incident from landing at all.

Infostealers are a SaaS Industry now

Across the engagements we worked this quarter, infostealer logs continue to show up in the timeline before the encryptor (or the leak site listing) lands. Industry breach investigation data published in April put the figure at roughly 73 percent of ransomware victims having an associated infostealer infection or credential leak in the prior year, with half of those infections inside the 95 days immediately before the attack. We see the same pattern in our caseload.

The Lumma operation, which was disrupted by international law enforcement last year, has resumed activity and is again one of the major credential collection operations in the world. Vidar 2.0, a full rewrite, has filled the gap left during the Lumma disruption; it is sold for around 300 dollars a month, supports multithreaded exfiltration, and bypasses the latest Chrome App Bound Encryption by direct memory injection.

The bigger story is VoidStealer, which appeared in March. It is the first stealer family in public knowledge to bypass Chrome App Bound Encryption without code injection or privilege escalation, using a debugger-based technique to extract the master key from process memory in plain text. Phemedrone, LummaC2, Meduza, StealC, Rhadamanthys and others have since added their own ABE bypass. The protective control that Google shipped in late 2024 has effectively been retired by the attacker side in under eighteen months.

The macOS side is no longer a curiosity. Atomic (AMOS) was distributed through a ClickFix-style campaign in April that used the Script Editor application to bypass the Terminal warning Apple introduced this year. A Microsoft-tracked ClickFix campaign on 6 May delivered AMOS alongside MacSync and SHub via fake macOS utility lures. There is also an active campaign that abuses ChatGPT shared conversations and paid Google ads to deliver AMOS with a persistent backdoor. The line that Macs do not need EDR protection should be retired this quarter.

Pricing for stealer subscriptions now starts at 30 dollars a month and reaches around 1,000 dollars for a lifetime licence. The average organisation now has somewhere around two and a half thousand corporate credentials per month being collected in stealer logs from devices it has no visibility over.

Twenty Two Seconds from Foothold to Handoff

Industry incident response telemetry published in late March recorded the most important operational statistic of the year. The median time between an initial access broker delivering a compromise and the secondary ransomware operator receiving and acting on that access is now around 22 seconds. Four years ago it was over eight hours. The handoff is no longer a market exchange that happens on a forum; it is an automated subprocess of the original infection.

22 secs

The median time from compromised access changing hands to ransomware deployment

4 years ago it was over 8 hours

What this means in practice is that the comfortable buffer between a low-severity initial alert and a full domain compromise no longer exists. By the time a SOC analyst opens a ticket on an initial malware detection the next morning, the secondary actor has had hours of hands-on keyboard access. Every initial access alert needs to be treated as a precursor to a ransomware event.

The initial access broker market itself has shifted in two directions this quarter. The average listing price on the major Russian language forums has risen to around 2,000 dollars, and premium tier listings on RAMP have moved six times higher than where they were a year ago. At the same time the volume of sub-1,000-dollar listings remains very high. There is now clear segmentation between low-cost generic access and high-value pre-validated full domain access with persistence.

The composition has not changed dramatically. VPN access without MFA, exposed RDP and domain user accounts continue to dominate the listings. UK manufacturing, hospitality, retail and education are all visible in recent listings.

We continue to investigate cases in which dormant or service accounts were used as the initial vector. In one engagement this quarter, we observed an access broker move a dormant account into an SSLVPN group, change the password and reset the MFA enrolment in a single sitting, all from a single compromised IT admin session, before selling the access on. That is the standard playbook now.

From our caseload

Patterns we have continued to see across Q2 engagements, with a few new arrivals.

1. **Help desk vishing** is now the most common identifiable initial vector across our retainer caseload. In several cases the call was placed not to the client directly but to their managed service provider's support desk. No malware. No exploit. Just a credible voice and a plausible story about a locked out user.

2. **Hypervisor level encryption** has become a default tactic. Akira in particular has driven the share of cases that involve direct encryption of ESXi virtual disks, bypass of in guest DLP and destruction of vCenter snapshots. We have seen Akira expand to Nutanix AHV environments from mid 2025 and the trend has continued through Q2. If your hypervisor management plane is reachable from your corporate network without segmentation, treat that as a critical exposure.

3. **Backup destruction** remains a defining feature of every ransomware engagement we worked this quarter. The Veeam vulnerabilities patched in March have moved Veeam squarely into the same priority bracket as edge VPN appliances for patching. We have observed authenticated low privilege accounts being used to execute remote code on Veeam backup servers, and seen the resulting destruction wipe immutable snapshots that were not in fact immutable because of misconfigured retention policies.

4. **Living off the land** remains the norm. PowerShell, WMIC, certutil, BITSAdmin and ntdsutil show up in essentially every case. We continue to see abuse of MSP RMM tools as the initial foothold; ConnectWise ScreenConnect has been particularly affected this quarter.

5. **Rclone** continues to be the exfiltration tool of choice, typically disguised as legitimate processes and dropped into directories like C:\PerfLogs or C:\Windows. FileZilla, AnyDesk and Sysinternals PSTools are usually somewhere in the same kill chain.

6. **Legal takedown** work has produced strong results again this quarter. Identifying the exfiltration destination (where the data was sent, where it is hosted, how it is being advertised) gives your legal team something operational to act on. We have had Q2 cases where rapid identification led to removal of stolen material before wide distribution.

7. **MSP compromise** continues to equal client compromise. We are still seeing precursor activity against MSP admin accounts months before the actual attack against the downstream client is executed.

8 Search engine poisoning to deliver fake installers continues to drive initial access. Any software download that arrives via search rather than via the vendor's own published link should now be treated as suspicious until proven otherwise.

9 Attackers targeting businesses in Business Email Compromise (BEC) cases are increasingly utilising custom-built toolkits to exploit compromised accounts. Among other capabilities, some of the toolkits can send direct API requests to Microsoft Office365 backend services, allowing attackers to quickly download messages en masse from compromised mailboxes. The end result of this shift in tactics for victims is that each email compromise can lead to a painful data exfiltration identification engagement.

Microsoft has built you new defences. Use them.

Q2 has been a busy quarter for Microsoft security capability releases, and several are directly relevant to the kinds of incidents we have just described.

The single biggest change is Automatic Attack Disruption in Defender XDR, which now includes automatic device isolation. This entered preview on 18 May. When high-confidence telemetry indicates a device is being used as an active attacker foothold, the device is automatically disconnected from the network at machine speed while remaining reachable by Defender for forensic collection. The action is time-limited, scoped to the incident, and can be released by an operator. Where this is enabled, we are seeing the attacker's command and control session severed within minutes of the first malicious sequence being detected. Given the 22-second handoff window described earlier, that is the difference between a contained incident and a catastrophic one.

The Contain User action reached general availability in March. It allows a suspicious identity to be temporarily contained at the endpoint layer without disabling the account in the identity provider, which preserves investigation telemetry while preventing further movement.

The Security Copilot Analyst Agent and the Alert Triage Agent reached general availability in Q2 and are now bundled into the Microsoft 365 E5 licence, with phased tenant rollout running from 20 April to 30 June. The Triage Agent is doing useful work on phishing alert volume.

Microsoft Agent 365, which reached general availability on 1 May, provides identity and governance for AI agents running on endpoints. As more organisations deploy copilots and agentic workflows, having those agents carry Entra identities and being governed by Defender, Purview and Intune is going to be load-bearing.

Entra Token Protection in Conditional Access cryptographically binds Primary Refresh Tokens to

the issuing device, which directly defeats the token replay pattern that the Code of Conduct, Kali365 and EvilTokens campaigns rely on. There are very few defensive controls available right now with this kind of impact for this kind of effort.

If you are on the Microsoft stack, these capabilities are not theoretical or far away. They are shipping now, several are included in licences you already pay for, and they directly counter the attack patterns we are investigating week to week.

CVEs we saw exploited across Q2 engagements

The following are the vulnerabilities our team observed actively exploited in Q2 engagements, or detected in client estates through proactive monitoring before they were used. Patch position on these **should not be left** to next month's window.

Veeam Backup and Replication. Multiple critical CVEs were disclosed on 12 March, versions **12** and **13**, including authenticated low-privilege RCE rated **CVSS 9.9**. Our investigators observed exploitation of these flaws in two Q2 ransomware engagements and detected unpatched Veeam instances in three further client environments through scheduled exposure scanning before any compromise occurred. Treat your backup servers as Tier Zero infrastructure and patch them on the same urgency window as your edge VPNs.

Cisco ASA and Firepower. CISA updated Emergency Directive 25-03 on 23 April after the ArcaneDoor activity using the **FIRESTARTER** implant. **CVE-2025-20333** (authenticated RCE as root) and **CVE-2025-20362** (unauthenticated URL endpoint access) chain together and the implant survives firmware upgrade because it lives in the FXOS base operating system. If you have not validated your ASA and Firepower estate against the directive, do it this week.

I**vanti Endpoint Manager Mobile. CVE-2026-1281** and **CVE-2026-1340** (both CVSS 9.8, unauthenticated RCE) were disclosed on 29 January 2026 and exploited within a day. Multiple European government bodies, including the European Commission, the Dutch Data Protection Authority and Finland's central government employer, have confirmed the compromise. New variants are still being added to KEV through Q2.

Fortinet FortiClient EMS, CVE-2026-35616, **CVSS 9.8**, actively exploited from 31 March.

Palo Alto Networks PAN-OS GlobalProtect, CVE-2026-0257. Authentication bypass via forged authentication override cookies. Earliest exploitation observed on 17 May, second wave 21 May. Added to **CISA KEV** on 29 May with federal remediation deadline 19 June.

ConnectWise ScreenConnect, **CVE-2026-3564** disclosed 20 March plus **CVE-2024-1708** added to KEV on 28 April. We observed unauthorised ScreenConnect sessions in two MSP-adjacent engagements this quarter.

Microsoft SharePoint Server, **CVE-2026-32201** actively exploited at the time of the April Patch Tuesday release.

Two supply chain campaigns also worth flagging. **Mini Shai-Hulud**, disclosed on 12 May, compromised over **160 npm** and **PyPI packages** including **TanStack**, **Mistral AI** and **Guardrails AI**, using a chain of three GitHub Actions vulnerabilities to inject self-propagating worm logic and a destructive daemon. **TrapDoor**, earliest observed activity 22 May, spread across **34 packages** and over **380 versions** across **npm**, **PyPI** and **Crates.io**. Both campaigns specifically target developer credentials, cloud provider tokens and Kubernetes service account tokens. If your build infrastructure consumes packages from these ecosystems, hunt for evidence of compromise on dependency installs from May onwards.

AI is now doing the attacker's work

We wrote in April about AI compressing the attack chain. In Q2 the public evidence has caught up with what we have been seeing.

Anthropic published in mid November 2025, and provided further detail across Q1 and Q2, what they described as the first reported AI orchestrated cyber espionage campaign. A Chinese state-aligned group used Claude Code agentically to conduct reconnaissance, exploit development, lateral movement and data exfiltration across approximately 30 organisations in technology, finance, chemicals and government. The AI performed between 80 and 90 percent of the work in each intrusion, with only four to six human decision points per engagement.

Microsoft, Google and OpenAI have all published disruption reports during Q2 documenting threat actor abuse of their own models. A North Korean cluster has been observed using developer platforms for staging and command and control. Multiple groups are using generative AI to draft phishing lures, translate content into target languages, summarise stolen data and generate or debug malware. Two malware families publicly tracked as PROMPTFLUX and PROMPTSTEAL query a hosted LLM at runtime to vary their behaviour and evade detection.

On voice, industry incident response telemetry now ranks voice phishing as the second most common initial infection vector at around 11 percent, with email phishing falling to around 6 percent. Contact centre deepfake fraud attempts have risen by around 1,300 percent year on

year. The cost of producing a convincing voice clone is now around three seconds of source audio for material that fools call centre verification systems.

On the email side, around 76 percent of phishing in 2024 contained at least one AI-driven polymorphic feature. AI-generated phishing achieves over 50 percent click-through, compared with around 12 percent for traditional templates.

Behind the campaigns sit a growing number of jailbroken or malicious commercial LLMs offered on Telegram and underground forums at prices starting around 50 dollars a week. The barrier to entry for AI-assisted offensive operations is now negligible.

The National Institute of Standards and Technology (NIST) maintains the AI Risk Management Framework (AI RMF), a voluntary framework developed to better manage risks to individuals, organisations, and society associated with artificial intelligence. First published in January 2023, it is built around four core functions, Govern, Map, Measure, and Manage, and is designed to incorporate trustworthiness considerations into the design, development, use, and evaluation of AI products, services, and systems. NIST has continued to expand it, releasing a Generative Artificial Intelligence Profile in July 2024 and, most recently, on April 7, 2026, a concept note for an AI RMF Profile on Trustworthy AI in Critical Infrastructure, which will guide critical infrastructure operators towards specific risk management practices to consider when engaging AI-enabled capabilities. For organisations adopting AI, it offers a structured, technology-neutral approach to identifying and mitigating risk without prescribing a one-size-fits-all standard

AI Risk Management Framework | NIST

Backups alone will not save you anymore

Recovery denial is now a category of threat in its own right, not a feature of one or two engagements. Industry data published this year shows that backup compromise is attempted in roughly 96 percent of ransomware attacks and succeeds in around 76 percent of those attempts. Where backups are compromised, the recovery cost runs around eight times higher (around three million dollars on average compared to around 375 thousand where backups remain intact) and the recovery window stretches from one week to four or more.

Hypervisor-focused encryption has risen from around 3 percent of cases in the first half of 2025 to around 25 percent in the second half. That is a 700 percent jump in under a year, driven

primarily by Akira (although Qilin and Black Basta affiliates have moved the same way). ESXi zero days disclosed earlier this year have been observed in active ransomware deployment, including some that were in the wild for more than a year before being patched.

The defensive position has not changed. If your backup accounts share credentials with your production domain, you are exposed. If your hypervisor management plane is reachable from the corporate network without segmentation, you are exposed. If your snapshots are not actually immutable (and a worrying number of installations we have inspected this year claim to have immutable backups but have allowed retention policies to be overridden by privileged accounts), you are exposed.

Test your restore capability against a scenario in which the attacker owns your domain admin and your backup admin. If the honest answer is that you cannot recover under those conditions, that is the gap to close this quarter, not next.

UK Snapshot

Some of what is happening in the UK regulatory and threat environment during Q2 will shape the rest of the year and beyond.

The **NCSC annual review**, published alongside **CYBERUK 2026** in Glasgow in late April, recorded around 204 nationally significant cyber attacks in the year to August 2025, up from 89 the year before. The NCSC CEO stated publicly that the majority of these are nation-state in origin and that the gap between threat and national defensive capacity is widening.

The Cyber Security and Resilience Bill, introduced to the Commons on 12 November 2025, completed Report Stage in Q1 and is now heading through the House of Lords. Royal Assent is expected this year, with phased implementation likely running into 2028. Penalties under the upper band will reach the higher of 17 million pounds or 4 percent of global turnover. UK businesses in scope of the expanded NIS regime should be planning for that timeline now.

On 12 May the ICO fined South Staffordshire Plc just under one million pounds over the August 2022 ransomware incident that exposed the data of more than 633,000 customers.

The UK cyber insurance market is the other moving piece. The reinsurance market reported a 32 percent reduction in risk-adjusted rate change on cyber aggregate excess of loss at the 1 January renewals, the most significant softening in the market's history. That softening is happening at the same time as US insurance premiums are rising sharply because of the M&S and Co-op events. Premium movement and incident frequency have decoupled this quarter. We expect that to correct rather than persist.

NCSC has published advisories during the quarter on Russian APT compromise of UK home and small office routers, Cisco Catalyst SD-WAN, F5 BIG-IP Access Policy Manager, Citrix NetScaler ADC and Gateway vulnerabilities, and heightened indirect cyber risk to UK organisations with Middle East regional exposure.

Cyber Essentials April 2026 Updates (CE Danzell)

Since the April 2026 update, Cyber Essentials has become more rigorous in both how organisations are assessed and how compliance is evidenced. While the five core technical controls remain unchanged, the scheme now places greater emphasis on demonstrating that security controls are **fully implemented** across the organisation's environment, rather than simply documented or planned.

Several of the changes focus on areas that are frequently exploited by attackers. **Multi-factor authentication (MFA)** is now an automatic fail requirement for cloud services where MFA is available, regardless of whether it is included as standard or requires an additional licence. Likewise, organisations must ensure that all **high-risk and critical security updates** for operating systems, applications, firewalls and routers are applied within **14 days of release**. Failure to meet these patching requirements can now result in an immediate assessment failure.

The update also introduces clearer scoping requirements. **Cloud services** can no longer be excluded from a Cyber Essentials assessment, and organisations operating a partial scope must now provide stronger justification for what has been excluded and how those systems are segregated from the certified environment. In addition, Cyber Essentials Plus assessments now require assessors to test a larger sample of devices. The introduction of the **"double sampling" rule** means that where issues are identified during testing, assessors may be required to expand the number of systems reviewed. This increases the likelihood that weaknesses, inconsistencies, or isolated compliance gaps will be identified during the assessment process and reinforces the importance of applying controls consistently across all in-scope devices.

For organisations pursuing Cyber Essentials Plus, the practical impact is clear: the gap between the self-assessment and the technical verification process has narrowed considerably. Organisations should expect **greater scrutiny** of the controls declared during the Cyber Essentials assessment, making it essential that any compliance claims can be evidenced in practice. Put simply, if you state that a control is in place, you should be **confident** that it will withstand independent verification during a Cyber Essentials Plus assessment.

What we expect through Q3

Calling the next quarter is harder than it has been in some time. With that caveat, here is what we think the second half of the summer looks like.

Cartel consolidation will continue to deepen. The Trinity of Chaos brand is unlikely to be the last public cartel announcement of the year, and we expect more aggressive cross-posting of victim data between affiliated leak sites. The traditional summer slowdown is unlikely to hold in 2026.

Pure extortion incidents will continue to grow as a share of the engagements we work. As payment rates continue to decline, the operators that can monetise without operational disruption have a structural advantage.

AI-driven vishing will scale. Synthetic voice cost has effectively reached zero, and the call centre and help desk attack surface has not moved. We expect one or more high-profile UK breaches in Q3 where the initial vector is conclusively a voice clone, and we expect that to drive a wave of help desk procedural reform across financial services and legal. Plan for that wave rather than react to it.

ClickFix and its variants (CrashFix, FileFix, and related "copy-and-paste" social engineering techniques) have continued to proliferate throughout the reporting period and remain one of the most prevalent initial-access vectors observed in the wild. The technique works by presenting victims with a fabricated error, CAPTCHA, or verification prompt - often on compromised or spoofed websites - that instructs them to manually copy a string, paste it into the Windows Run dialog (or a terminal), and execute it, thereby running attacker-supplied PowerShell or shell commands under their own user context.

Because the victim performs the execution themselves, the approach neatly sidesteps many browser download warnings, mark-of-the-web protections, and automated email/attachment filtering. We have observed ClickFix-style lures being used to deliver a broad range of payloads, including infostealers (Lumma, StealC),

RATs, and loaders, and the tradecraft has been adopted by both opportunistic criminal operators and more capable threat groups.

The variants differ mainly in delivery pretext - CrashFix leaning on fake browser/application crash dialogs and FileFix abusing the file-explorer address bar - but all share the same core weakness: they exploit the user as the execution mechanism rather than a software vulnerability, which makes user awareness training and restriction of the Run dialog / scripting hosts the most effective countermeasures.

The Microsoft Defender Automatic Attack Disruption preview is likely to move toward general availability in Q3. Organisations that piloted it during the preview will be materially better placed than those that did not.

Identity infrastructure CVEs will continue to land. The Microsoft Entra Actor Token disclosure earlier this year has set the tone. We expect at least one further serious identity provider vulnerability in Q3 and rapid weaponisation when it happens.

Iran state-aligned cyber activity will continue to drive baseline risk for any UK organisation with regional exposure or supply chain dependency. The campaign that began in February with retaliatory operations has not slowed. The UK is not the primary target but the indirect risk to UK firms with US parents or US-dependent supply chains is material.

The UK Cyber Security and Resilience Bill is likely to receive Royal Assent during Q3. Get your boards across what in scope means for your sector now.

The next round of ransom payment rate data is expected to be published during Q3. Based on the late 2025 trajectory, we anticipate the reported payment rate will fall to around 20 percent. If you are in a sector where payment is still the assumed default, that conversation needs to happen at board level before the next incident, not during it.

CYFOR SECURE
CYBER SECURITY

The Bottom Line

What every firm handling sensitive data should have in place during the second half of 2026.

- **Disable OAuth device code flow** in Entra Conditional Access unless your tenant explicitly needs it. The single biggest reduction in Q2 phishing exposure available to most organisations sits behind one toggle.
- **Enable Entra Token Protection** in Conditional Access. It cryptographically binds tokens to the issuing device and directly defeats the replay pattern used by Kali365, EvilTokens, and the Code of Conduct campaign.
- **Move privileged accounts** onto phishing-resistant MFA (FIDO2 or Windows Hello for Business). Push notification and SMS based MFA should be considered legacy controls.
- **Harden your help desk.** Out-of-band callback verification on every password reset and MFA re-enrolment for privileged accounts. Manager approval step for high-risk requests. Train staff to assume voice cloning.
- **Audit remote access accounts.** Dormant accounts, former staff, shared credentials. If a VPN account has not been used in 90 days, disable it today.
- **Patch the priority list now.** Veeam, Cisco ASA and Firepower, Ivanti EPMM, Fortinet FortiClient EMS, Palo Alto GlobalProtect, ConnectWise ScreenConnect. These are not future risks. They are current ones with confirmed exploitation.
- **Test backup resilience under domain compromise.** Can you restore if the attacker owns your domain admin and your backup admin? Are backup credentials independent of production AD? Are snapshots actually immutable or just configured to be? If you cannot answer yes to all three, recovery denial will work against you.
- **Segment your hypervisor management plane.** ESXi, vCenter, and Hyper-V management interfaces should not be reachable from your standard corporate user network.

CYFOR SECURE
CYBER SECURITY

The Bottom Line Cont.

What every firm handling sensitive data should have in place during the second half of 2026.

- Monitor for credential and infrastructure exposure outside the perimeter. Your domain almost certainly appears in stealer logs from devices you do not manage; your network may already be advertised by an access broker. If you do not have continuous visibility over this, you are operating without one of the most useful early warning signals available today.
- Review your MSP's access. What accounts do they hold against your environment? What MFA protects them. If your MSP is compromised, you are compromised. Ask for that evidence in writing this quarter.
- Adopt the Microsoft defensive stack you already pay for. Automatic Attack Disruption, Contain User, Security Copilot Triage Agent, Token Protection. These are not future capabilities.
- Plan for AI-driven social engineering. Voice cloning, AI-generated phishing, synthetic identities. Your awareness training needs to reflect what an attack looks like today, not what it looked like in 2024.

CYFOR Secure provides digital forensics, incident response, and expert witness services across the UK. If you have experienced a breach, suspect one, or want to understand your exposure - get in touch.

0330 135 5756 contact@cyforsecure.co.uk cyforsecure.co.uk