ultraviolet

QUARTERLY REPORT 2026

Q2 Threat Review

Threat Actors, Exploited Vulnerabilities & Social Engineering

PREPARED BY
UltraViolet Cyber TIDE Team
tide@uvcyber.com

Table of Contents

Executive Snapshot

The second quarter of 2026 reinforced a pattern that has been building for several consecutive quarters: threat actors are becoming more operationally disciplined, attack surfaces are expanding faster than most organizations can secure them, and the human element remains the most reliable entry point for adversaries at every level of sophistication.

47%	**11%**	**$25M+**
of initial access incidents traced to the ClickFix technique	of confirmed initial access from vishing — now ahead of email at 6%	lost in a single deepfake video impersonation incident

On the ransomware front, the dominant groups of Q2 operated less like opportunistic criminals and more like structured enterprises. Qilin, Cl0p, Akira, and DragonForce each demonstrated refined targeting strategies, with a continued shift toward data-theft-based extortion over traditional encryption. DragonForce's cartel-style affiliate model and Cl0p's pivot to direct executive-level pressure illustrate a broader evolution in which ransomware operators are optimizing for leverage and business disruption rather than volume. Meanwhile, Salt Typhoon and affiliated Chinese state-sponsored actors continued pre-positioning operations across North American telecommunications and critical infrastructure, blurring the line between cybercrime and geopolitical strategy.

The vulnerability landscape this quarter was shaped by both legacy exposure and newly disclosed flaws. The Cisco Catalyst SD-WAN authentication bypass (CVE-2026-20127) dominated early Q2 response efforts, with CISA issuing Emergency Directive 26-03 and a cascade of related SD-WAN vulnerabilities surfacing in the weeks that followed. The cPanel/WHM authentication bypass saw mass exploitation by the "Sorry" ransomware group, while Fortinet and BeyondTrust flaws continued to give attackers footholds across edge-facing infrastructure. Late in the quarter, the public disclosure of YellowKey introduced a BitLocker bypass affecting Windows 11 and Windows Server 2022/2025 that requires no credentials or recovery key, only physical access and a crafted USB drive. Released without a CVE or patch one day after May's Patch Tuesday, YellowKey is a pointed reminder that trusted native components like the Windows Recovery Environment can become part of the attack surface in ways that traditional detection models are not built to catch.

Social engineering continued its trajectory as the most reliable initial access method across the threat landscape. AI-generated phishing, which surged dramatically in late 2025, held at elevated levels throughout Q2, with AI-crafted lures producing substantially higher engagement rates than their human-written counterparts. Perhaps more notable is the displacement of email by voice phishing as the leading social engineering vector in confirmed incidents. Attackers are calling help desks, cloning executive voices with seconds of publicly available audio, and conducting real-time impersonation that leaves no malware trace. Pretexting-driven business email compromise has nearly doubled, now representing the majority of all social engineering incidents, with attackers engaging finance teams in multi-turn conversations rather than relying on single-click lures. North Korea's operationalization of deepfake IT worker infiltration added a new dimension to the threat, embedding state-sponsored operatives directly into Western payrolls for espionage and revenue generation.

Taken together, the Q2 2026 landscape points to a consistent theme: the gap between attacker capability and defender response time is widening. Exploitation windows are shrinking to days, not weeks. Social engineering is outpacing awareness training. Trusted components are being weaponized. The organizations best positioned to weather this environment are those treating security as a continuous operational discipline rather than a periodic compliance exercise, with investments concentrated in identity verification, zero-trust architecture, rapid patching workflows, and resilience against the attacks that bypass technical controls entirely.

Throughout this report are the Top 5 Threat Actor Groups, Top 5 Vulnerabilities, and Top 5 Social Engineering techniques seens over Q2 2026. A special section on ClickFix style attacks has been included as well.

SECTION 2

Top 5 Threat Actor Groups

02

The ransomware ecosystem matured into structured enterprises, and state-sponsored actors kept pre-positioning inside critical infrastructure.

1 Qilin

RANSOMWARE-AS-A-SERVICE

SECTORS Healthcare, construction, IT & professional services	**REGIONS** US, Europe, APAC
MODEL RaaS — affiliate network	**NOTABLE** Rust variants, ESXi targeting

Qilin has established itself as one of the most operationally active ransomware groups of Q2 2026, maintaining a consistent pace of victim postings on leak sites throughout the quarter. The group has demonstrated a preference for targeting healthcare, construction, IT services, and professional services organizations, with a geographic reach spanning the United States, Europe, and the Asia-Pacific region. Their sustained daily operations reflect a mature and well-resourced affiliate network capable of running parallel campaigns across multiple sectors simultaneously.

A defining characteristic of Qilin's technical evolution has been the adoption of Rust-based ransomware variants designed to complicate reverse engineering and hinder analysis by defenders. These variants are capable of targeting both Windows and Linux environments, including VMware ESXi virtual machines, which broadens the group's ability to disrupt enterprise infrastructure at scale. Their tooling updates signal a commitment to cross-platform targeting that aligns with the broader industry trend of ransomware operators pursuing hypervisor-level encryption for maximum operational impact.

Qilin operates under a Ransomware-as-a-Service model, recruiting affiliates through underground forums and providing them with customizable payloads, negotiation infrastructure, and leak site hosting. This decentralized execution model allows the group to scale rapidly while insulating core operators from direct exposure. Their consistent output across Q2 suggests that affiliate recruitment and retention remain strong, and organizations in their preferred target verticals should treat Qilin as a persistent and credible threat.

Action Items

- ☐ Prioritize patching of internet-facing infrastructure, particularly VPN appliances, remote access tools, and virtualization platforms, as these represent common initial access vectors for Qilin affiliates.
- ☐ Implement network segmentation to limit lateral movement, with particular attention to isolating VMware ESXi hosts and management interfaces from general user networks.
- ☐ Ensure offline, immutable backup copies exist for critical systems and test restoration procedures regularly, as Qilin's dual extortion model targets both data availability and confidentiality.
- ☐ Deploy endpoint detection and response tooling with behavioral detection capabilities tuned for Rust-based payloads and cross-platform encryption activity.

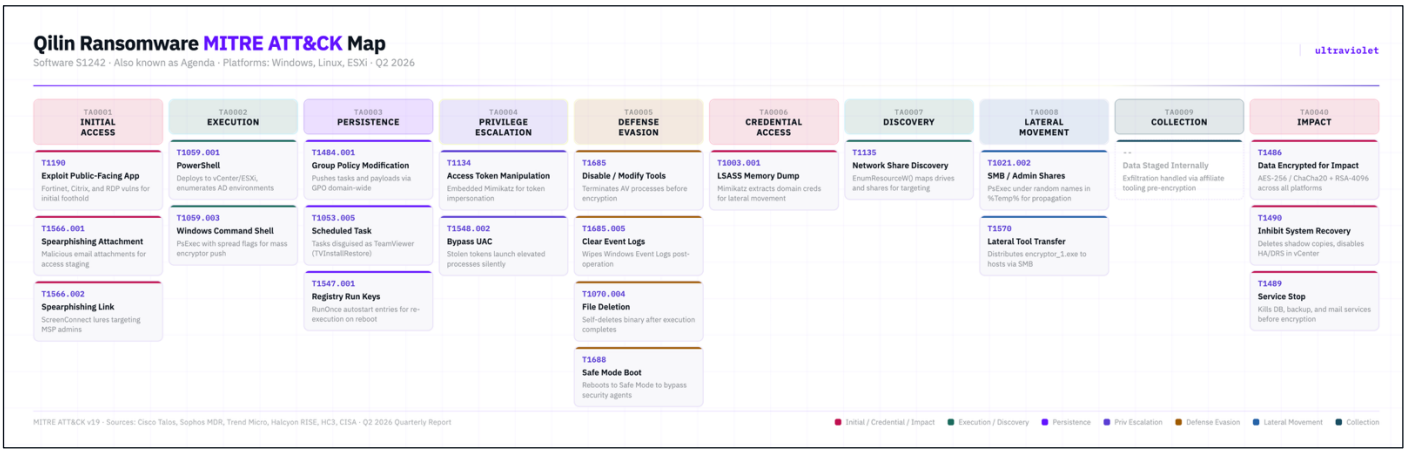

2 Cl0p

DATA-THEFT

SECTORS Cross-industry via supply chain	**REGIONS** Global
MODEL Zero-day MFT exploitation	**NOTABLE** Executive-level pressure

Cl0p continued to distinguish itself in Q2 2026 through a strategic pivot away from traditional ransomware encryption and toward pure data-theft extortion. Rather than deploying ransomware lockers to encrypt victim systems, Cl0p focuses on exfiltrating sensitive data early in the intrusion lifecycle and then pressuring victims through the threat of public exposure on their leak site. This approach reduces the group's detection footprint, shortens dwell time requirements, and exploits reputational damage as the primary leverage mechanism against targeted organizations.

A notable tactical evolution this quarter has been Cl0p's increasing reliance on supply chain exploitation and targeting of managed file transfer platforms and enterprise SaaS applications. By compromising widely used third-party software, Cl0p is able to reach dozens or even hundreds of downstream victims from a single intrusion. This approach to systemic exploitation allows the group to operate at a scale that would be impossible through traditional one-at-a-time targeting, and it takes advantage of the trust relationships that organizations place in their software vendors and service providers.

Action Items

- ☐ Audit and harden all managed file transfer solutions and enterprise SaaS platforms, ensuring they are fully patched and configured according to vendor security guidance.
- ☐ Implement data loss prevention controls and monitor for anomalous large-scale data exfiltration, particularly from file sharing and collaboration platforms.
- ☐ Establish executive-specific communication protocols so that leadership teams are prepared for direct extortion contact and know to route such communications through legal and incident response channels.
- ☐ Conduct supply chain risk assessments for critical third-party vendors, with emphasis on understanding their patch management practices and breach notification commitments.

Cl0p has also refined its extortion workflow to include direct outreach to executive leadership at victim organizations, bypassing IT and security teams entirely. This shift toward executive-level pressure campaigns reflects a sophisticated understanding of corporate decision-making dynamics and increases the likelihood of payment by creating urgency at the highest levels of the organization. The combination of supply chain access, data-only extortion, and executive targeting makes Cl0p one of the most strategically dangerous groups operating in the current landscape.

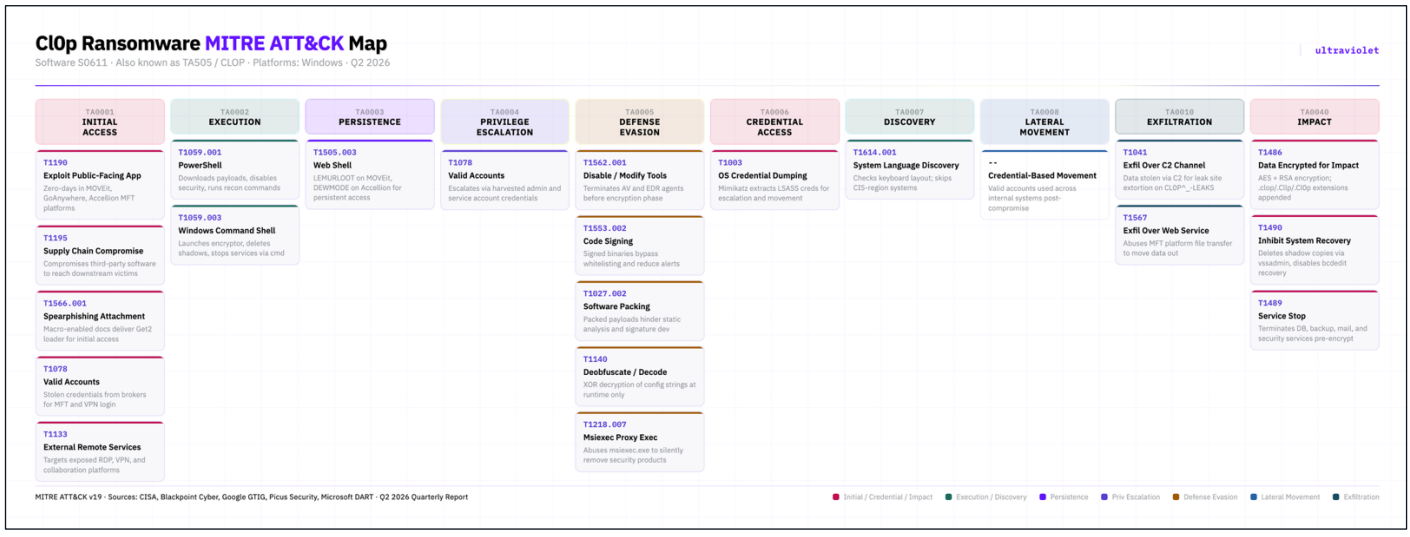

3 Akira

DOUBLE EXTORTION

SECTORS Manufacturing, professional services, healthcare, education	**REGIONS** Europe & DACH
MODEL Double extortion	**NOTABLE** C++ & Rust ESXi variants

Akira has maintained a steady operational cadence throughout Q2 2026, continuing to target manufacturing, professional services, healthcare, and education sectors with a particular concentration in Europe and the DACH region. The group's targeting strategy is deliberate rather than opportunistic, focusing on industries where operational disruption creates cascading economic, regulatory, and reputational consequences. This calculated approach to victim selection reflects a mature understanding of which organizations are most likely to pay ransoms and which sectors face the steepest costs from prolonged downtime.

Action Items

- ☐ Harden remote access infrastructure, including VPN gateways and remote desktop services, as Akira affiliates frequently exploit unpatched or misconfigured remote access as an initial entry point.
- ☐ Enforce multi-factor authentication across all externally facing services and privileged internal accounts, with preference for phishing-resistant methods such as FIDO2 or hardware tokens.
- ☐ Develop and regularly test incident response playbooks that account for dual extortion scenarios, including both system recovery and data breach notification obligations.
- ☐ Engage with sector-specific threat intelligence sharing organizations to stay current on Akira's evolving tactics and targeting patterns within your industry vertical.

Technically, Akira operates a double extortion model in which victim data is both encrypted and exfiltrated. The group demands payments ranging from hundreds of thousands to several million dollars depending on the size and perceived ability to pay of the target organization. Akira's tooling has continued to evolve, with recent variants built in C++ for improved stability and performance, as well as Rust-based Linux variants designed specifically for targeting VMware ESXi environments. This dual-platform capability allows Akira to threaten both traditional Windows-based infrastructure and virtualized environments in a single engagement.

Akira's geographic pivot toward European industrial hubs is not random. These regions function as logistical and manufacturing centers where even brief periods of downtime carry immediate financial impact across interconnected supply chains. By targeting organizations embedded in these networks, Akira amplifies its leverage beyond the individual victim and creates pressure that extends to partners, customers, and regulators. Organizations operating in these sectors and regions should consider themselves within Akira's primary targeting profile.

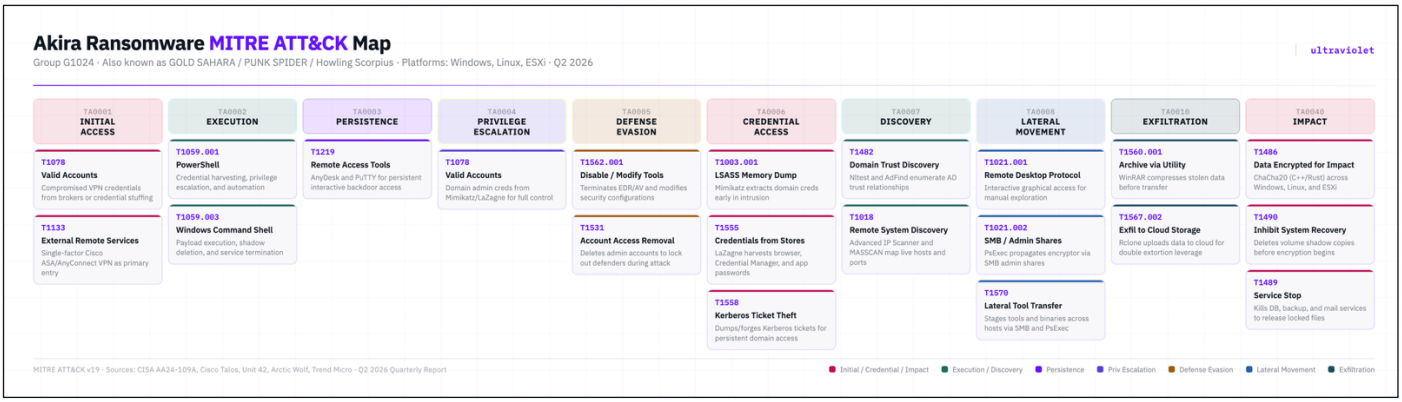

4 DragonForce

CARTEL-STYLE RAAS

SECTORS Manufacturing, retail, financial services, healthcare	**REGIONS** Global
MODEL Cartel-style affiliate RaaS	**NOTABLE** LockBit 3.0 & Conti V3-derived

DragonForce has rapidly scaled its operations over the past several quarters by adopting a cartel-style Ransomware-as-a-Service model that centralizes key functions while distributing attack execution across a broad affiliate network. The group manages core infrastructure including leak sites, negotiation channels, and payment processing, while affiliates handle intrusion operations independently. This division of labor allows DragonForce to increase its attack volume and geographic reach without proportionally increasing the operational burden on its core team, and it makes law enforcement disruption significantly more difficult.

The group's victim profile spans manufacturing, retail, financial services, and healthcare, with a focus on organizations that hold high volumes of personally identifiable information or operate in sectors where business interruption costs are severe. DragonForce's attacks frequently result in multi-day operational outages and exposure of millions of customer records, driving ransom demands into the multi-million-dollar range. Their willingness to publish stolen data aggressively and quickly on their leak site adds urgency to victim negotiations and reduces the window for organizations to mount an effective response.

DragonForce's rapid growth trajectory makes it a group to watch closely in the second half of 2026. The cartel model is proving effective at attracting skilled affiliates, and the group's centralized infrastructure provides a level of operational consistency that many newer ransomware operations lack. As DragonForce continues to mature, organizations should expect increasingly targeted and well-resourced campaigns, particularly in sectors with high data sensitivity and low tolerance for downtime.

Action Items

- ☐ Conduct regular external attack surface assessments to identify and remediate exposed services, default credentials, and unpatched systems that affiliates may target for initial access.
- ☐ Implement robust data classification and access control policies to limit the volume of sensitive data accessible from any single compromised account or system.
- ☐ Prepare public communications and regulatory notification templates in advance so that the organization can respond rapidly in the event of data exposure on a leak site.
- ☐ Evaluate cyber insurance coverage to ensure it accounts for dual extortion scenarios, including both business interruption costs and third-party liability from data exposure.

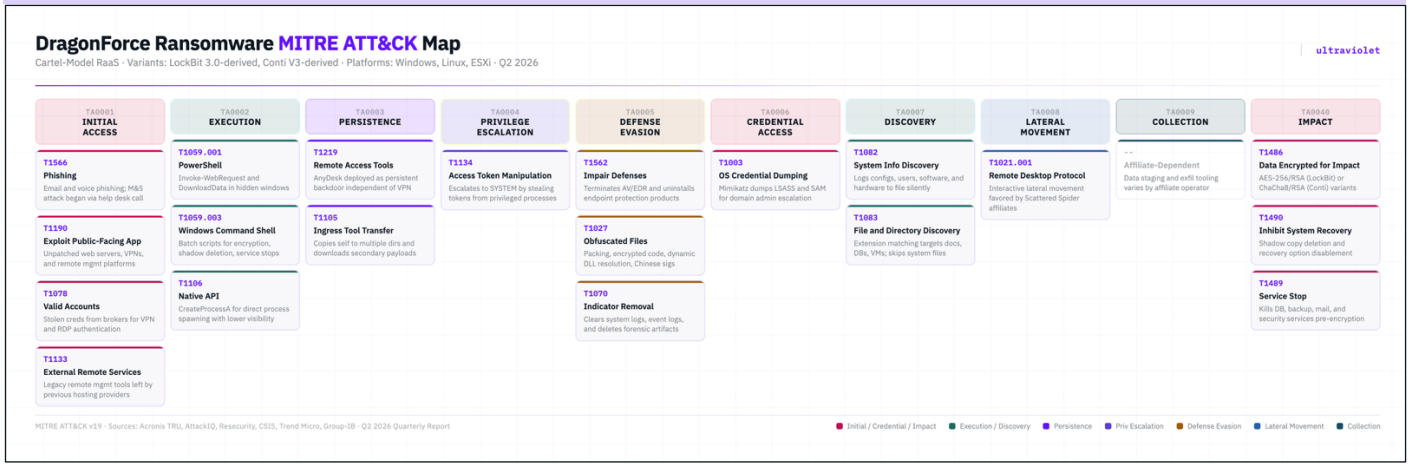

5 Salt Typhoon

STATE-SPONSORED

SECTORS Telecom, government, IT services, defense-adjacent	**REGIONS** North America
MODEL Living-off-the-land APT	**NOTABLE** Long-term pre-positioning

Salt Typhoon represents one of the most significant state-sponsored threats of Q2 2026, operating as part of a broader ecosystem of Chinese advanced persistent threat groups focused on long-term strategic positioning within Western critical infrastructure. The group has been primarily targeting North American telecommunications providers, government networks, commercial IT services, and defense-adjacent organizations. Unlike financially motivated ransomware operators, Salt Typhoon's objective is not immediate monetization but rather the establishment of persistent access that can be leveraged for espionage, intelligence collection, or disruptive operations at a time of geopolitical advantage.

A defining feature of Salt Typhoon's operational approach is the use of living-off-the-land techniques, in which the group relies on legitimate system tools and trusted native processes rather than custom malware to move through compromised environments. This approach makes detection exceptionally difficult, as the activity blends seamlessly with normal administrative operations. The group has also demonstrated a preference for targeting edge-facing network infrastructure, including routers, firewalls, and SD-WAN controllers, where visibility and monitoring are often weaker than on endpoints and servers within the corporate network.

The strategic implications of Salt Typhoon's activities extend well beyond individual victim organizations. By pre-positioning within telecommunications and IT service providers, the group gains potential access to the communications and data of downstream customers, government agencies, and critical infrastructure operators. This makes Salt Typhoon's campaigns a national security concern rather than a conventional cybersecurity incident. Organizations in the telecommunications, government, and critical infrastructure sectors should operate under the assumption that state-sponsored actors are actively seeking footholds in their environments.

Action Items

- ☐ Conduct thorough reviews of edge-facing network infrastructure, including routers, firewalls, VPN concentrators, and SD-WAN controllers, for signs of unauthorized configuration changes, unexpected firmware versions, or anomalous administrative access.
- ☐ Implement enhanced logging and monitoring for living-off-the-land activity, including unusual use of native system administration tools such as PowerShell, WMI, and command-line utilities by non-administrative accounts.
- ☐ Participate in government and sector-specific threat intelligence sharing programs such as those coordinated by CISA, NSA, and relevant ISACs to stay informed on the latest indicators of compromise and tactical guidance related to state-sponsored threats.

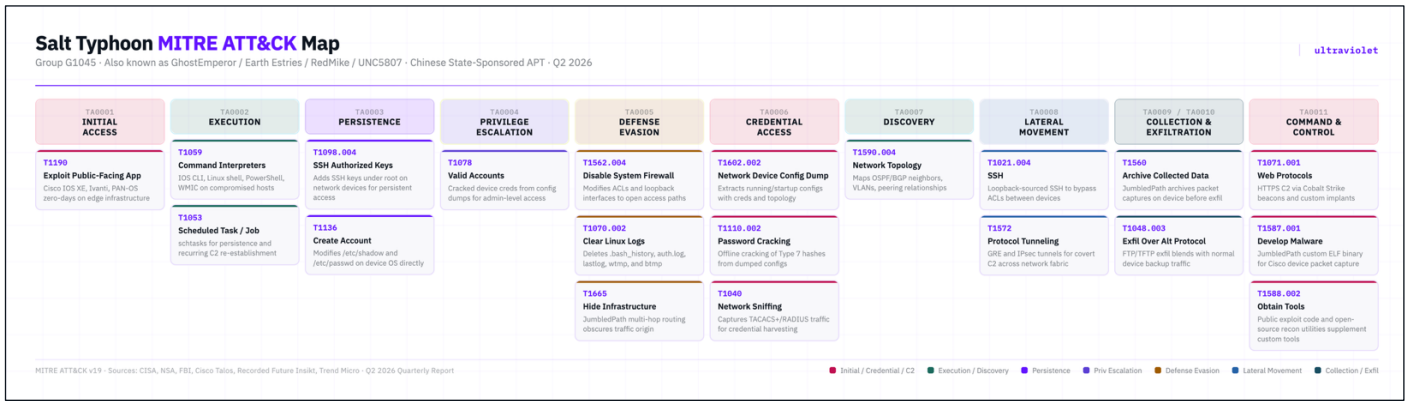

SECTION 3

05 Top 5 Exploited Vulnerabilities

Edge-facing infrastructure was the preferred entry point — chained, exploited for years before disclosure, and patched too slowly.

CVE-2026-20127 10.0 CRITICAL

1 Cisco Catalyst SD-WAN Authentication Bypass

CISCO CATALYST SD-WAN CONTROLLER & MANAGER

CVE-2026-20127 was the single most consequential vulnerability of Q2 2026, affecting Cisco Catalyst SD-WAN Controller and Cisco Catalyst SD-WAN Manager deployments worldwide. The flaw exists in the peering authentication mechanism of affected systems and allows an unauthenticated remote attacker to bypass authentication entirely and obtain administrative privileges by sending crafted requests. Once inside, an attacker can access the NETCONF service, manipulate network configuration across the entire SD-WAN fabric, inject rogue peers into the management plane, and begin establishing long-term persistence. The vulnerability carries a maximum CVSS score of 10.0 and requires no user interaction or prior authentication to exploit.

What makes this vulnerability particularly alarming is the timeline of exploitation. Although Cisco disclosed the flaw in late February 2026, Cisco Talos attributed the activity to a highly sophisticated threat actor tracked as UAT-8616 and assessed with high confidence that exploitation had been occurring since at least 2023. The attackers were observed chaining CVE-2026-20127 with CVE-2022-20775, a previously known privilege escalation bug in Cisco SD-WAN CLI, to downgrade compromised devices to vulnerable software versions, escalate to root-level access, and then restore the original software version to cover their tracks. This level of operational discipline suggests a well-resourced, likely state-aligned actor with deep familiarity with Cisco's SD-WAN architecture.

The response to this vulnerability was immediate and sweeping. CISA issued Emergency Directive 26-03 on the same day as disclosure, requiring all federal civilian agencies to inventory SD-WAN systems, collect forensic artifacts, apply patches, and hunt for compromise within an accelerated two-day window. A joint advisory was co-authored by CISA, NSA, and international partners from Australia, Canada, New Zealand, and the United Kingdom, accompanied by a detailed threat hunt guide. In the weeks that followed, three additional Cisco SD-WAN vulnerabilities were confirmed as actively exploited: CVE-2026-20128, CVE-2026-20122, and CVE-2026-20133. Organizations running Cisco Catalyst SD-WAN in any capacity should treat this as one of the highest-priority remediation efforts of 2026.

Action Items

- ☐ Immediately verify that all Cisco Catalyst SD-WAN Controller and Manager instances have been upgraded to the fixed software versions specified in Cisco's advisory, regardless of whether exploitation has been detected.
- ☐ Review SD-WAN logs for anomalous control connection peering events, unexpected reboots, version downgrade activity, and the creation of unfamiliar local user accounts, following the detection guidance published in the joint CISA/NSA threat hunt guide.
- ☐ Restrict network exposure of SD-WAN control plane interfaces, ensuring that NETCONF (TCP/830) and management interfaces are not accessible from the public internet and are protected by firewall rules and strong authentication.
- ☐ Conduct a full forensic assessment of SD-WAN infrastructure even if patches have already been applied, as attackers may have established persistence mechanisms prior to remediation.

CVE–2026-45585 6.8 - MEDIUM

2 YellowKey – BitLocker Bypass via WinRE

WINDOWS 11, WINDOWS SERVER 2022 / 2025

YellowKey is a BitLocker bypass vulnerability disclosed on May 13, 2026, one day after Microsoft's May Patch Tuesday, by a security researcher known as Chaotic Eclipse. The exploit targets the Windows Recovery Environment and allows an attacker with physical access to gain unrestricted shell access to a BitLocker-protected volume without any recovery key, credentials, or authentication. The attack is executed by placing a specially crafted "FsTx" folder on a USB drive or within the EFI partition, booting into WinRE using a specific key sequence, and triggering vulnerable recovery behavior that exposes the encrypted volume. The proof-of-concept was independently reproduced by security researcher Will Dormann shortly after publication, confirming the behavior and raising additional concerns about cross-volume file manipulation within the recovery environment.

The vulnerability affects Windows 11 and Windows Server 2022 and 2025. Windows 10 is not affected. The researcher noted that the component responsible for the vulnerable behavior exists only within the WinRE image and is not present in the same form in a standard Windows installation, despite a component with the same name existing in both environments. This discrepancy, combined with the fact that the flaw is limited to newer operating system versions, has prompted speculation about whether the behavior was intentional, though no evidence of a deliberate backdoor has been confirmed. As of this writing, Microsoft has not issued a CVE or a patch for YellowKey.

While the physical access requirement limits the scalability of this exploit, it represents a meaningful threat to organizations that rely on BitLocker as their primary data-at-rest protection for mobile devices, field laptops, shared workstations, and kiosk systems. Any device that can be physically accessed by an unauthorized party, whether through theft, insider access, or weak physical security controls, is potentially vulnerable. YellowKey was released alongside GreenPlasma, a related privilege escalation exploit targeting ctfmon.exe, by the same researcher who had previously disclosed the BlueHammer, RedSun, and UnDefend vulnerabilities earlier in 2026. BlueHammer was eventually patched by Microsoft only after active exploitation was observed in the wild.

Action Items

- ☐ Enforce BitLocker configurations that require a pre-boot PIN in addition to TPM protection, as the published exploit targets TPM-only configurations. Monitor Microsoft guidance closely, as the researcher has suggested the vulnerability may extend to TPM+PIN setups.
- ☐ Restrict removable media usage through group policy or endpoint management tooling to prevent unauthorized USB devices from introducing the crafted files required for exploitation.
- ☐ Lock BIOS/UEFI boot order and enable firmware passwords on all managed devices to prevent unauthorized boot into the Windows Recovery Environment.
- ☐ Review physical security controls for devices most at risk, including field laptops, shared workstations, kiosk systems, and any hardware that regularly leaves direct organizational control. Cable locks, locked storage, and check-in/check-out procedures reduce the window of physical access.
- ☐ Monitor for unexpected WinRE usage across the device fleet, as most end users have no legitimate reason to be in the recovery environment, and flag any such events for investigation.

CVE-2026-41940 9.8 - CRITICAL

3 cPanel/WHM Authentication Bypass

CPANEL & WEB HOST MANAGER

CVE-2026-41940 is a critical authentication bypass vulnerability in cPanel and Web Host Manager, two of the most widely deployed web hosting control panels in the world. The flaw allows an unauthenticated attacker to bypass login mechanisms and gain administrative access to the hosting server, from which they can modify configurations, access hosted websites and databases, and deploy malicious payloads across all accounts on the server. Given that a single cPanel instance may host dozens or hundreds of websites, the blast radius of a successful exploitation is substantial.

The vulnerability has been actively and aggressively exploited since February 2026 by the "Sorry" ransomware group, which has used it as the primary entry point for a campaign targeting shared hosting environments. After gaining access through the authentication bypass, attackers deploy a Linux-based encryptor that encrypts files across the server, affecting every hosted site and customer simultaneously. This approach allows the ransomware operators to impact a large number of victims from a single compromise, creating pressure on hosting providers to pay ransoms quickly in order to restore service for their entire customer base. The campaign has been particularly effective against smaller hosting providers and managed service providers that may lack the security maturity to detect and respond to the intrusion before encryption is complete.

The mass exploitation of CVE-2026-41940 underscores the outsized risk that shared infrastructure vulnerabilities pose to the broader ecosystem. A single unpatched cPanel server can serve as the entry point for compromising hundreds of downstream websites, customer databases, and email accounts. Hosting providers and managed service providers bear a disproportionate responsibility for patching this class of vulnerability quickly, as their customers have no ability to mitigate the flaw themselves. cPanel released patches for CVE-2026-41940 alongside two additional vulnerabilities (CVE-2026-29201, CVE-2026-29202, and CVE-2026-29203), all of which should be applied together.

Action Items

- ☐ Apply all available cPanel and WHM patches immediately, including fixes for CVE-2026-41940 and the related vulnerabilities disclosed in the same advisory cycle.
- ☐ Audit cPanel server access logs for signs of unauthorized administrative sessions, unexpected account creation, configuration changes, or file modifications dating back to February 2026.
- ☐ Implement network-level access restrictions for cPanel and WHM administrative interfaces, limiting access to known management IP addresses and requiring VPN connectivity for remote administration.
- ☐ If operating as a hosting provider or managed service provider, notify downstream customers of the vulnerability and any remediation actions taken, and conduct a thorough review to confirm that no unauthorized modifications were made to hosted sites or databases.

CVE-2026-21643 9.8 - CRITICAL

4 Fortinet SQL Injection

FORTIOS / FORTIGATE & RELATED

CVE-2026-21643 is a SQL injection vulnerability affecting Fortinet products that was added to CISA's Known Exploited Vulnerabilities catalog in April 2026 based on confirmed evidence of active exploitation. SQL injection flaws in network security appliances are particularly dangerous because these devices sit at the perimeter of the network and are, by design, reachable from the internet. Successful exploitation can allow an attacker to read or modify data within the appliance's management database, potentially extracting credentials, configuration data, and session tokens that enable deeper access into the protected network.

Fortinet edge devices have been among the most consistently targeted products in the vulnerability landscape for several consecutive years, and Q2 2026 continued that trend. Both ransomware affiliates and advanced persistent threat groups have demonstrated a strong preference for exploiting Fortinet appliances as initial access vectors, leveraging known vulnerabilities in FortiOS, FortiGate, and related products to establish footholds before moving laterally into internal environments. The recurring nature of these exploits reflects both the widespread deployment of Fortinet products and the difficulty many organizations face in maintaining timely patching on appliances that are considered critical infrastructure and are often difficult to take offline for maintenance.

The exploitation of CVE-2026-21643 should be understood in the context of a broader pattern in which edge-facing network infrastructure has become the preferred entry point for sophisticated attackers. Firewalls, VPN concentrators, and network management platforms represent high-value targets precisely because they are trusted, privileged, and externally accessible. Organizations that delay patching these devices, whether due to change management friction, availability concerns, or simple oversight, are accepting a level of risk that is increasingly untenable given the speed at which exploitation follows disclosure.

Action Items

- ☐ Apply the vendor-provided patch for CVE-2026-21643 immediately across all affected Fortinet deployments, prioritizing internet-facing appliances.
- ☐ Review Fortinet appliance logs for indicators of SQL injection attempts, unexpected administrative access, configuration exports, or credential harvesting activity.
- ☐ Implement a dedicated patch management process for edge-facing security appliances that operates on an accelerated timeline separate from standard infrastructure patching cycles, with the goal of applying critical patches within days rather than weeks of release.
- ☐ Reduce the attack surface of Fortinet management interfaces by restricting administrative access to internal management networks, disabling unnecessary services, and enforcing multi-factor authentication for all administrative sessions.

CVE-2026-1731 9.8 - CRITICAL

5 BeyondTrust Command Injection

BEYONDTRUST REMOTE / PRIVILEGED ACCESS MGMT

CVE-2026-1731 is a command injection vulnerability in BeyondTrust remote access and privileged access management software that allows an attacker to execute arbitrary commands on the underlying system without prior authentication. Remote access and privileged access management tools occupy a uniquely sensitive position in enterprise environments because they are specifically designed to broker access to critical systems, store privileged credentials, and facilitate administrative sessions. A vulnerability in this class of software does not simply compromise one system; it potentially compromises the keys to every system the tool is configured to manage.

The vulnerability was observed during Q1 and Q2 2026 being actively leveraged in advanced persistent threat campaigns involving multiple command-and-control frameworks, including Metasploit, Sliver, and Havoc. Attackers exploited CVE-2026-1731 alongside other authentication bypass vulnerabilities, such as CVE-2024-12356, to gain initial access to target environments and then used the compromised BeyondTrust infrastructure as a pivot point for deeper network penetration. The pairing of this vulnerability with established C2 frameworks indicates that it has been integrated into the standard operational playbooks of multiple threat actor groups, including those with likely state-aligned objectives.

The exploitation of privileged access management infrastructure represents one of the highest-impact attack patterns in the current landscape. These tools are trusted by design, whitelisted by security controls, and granted broad network access as a function of their intended purpose. When compromised, they provide attackers with a legitimate and often unmonitored pathway into the most sensitive parts of an organization's environment. The fact that CVE-2026-1731 requires no authentication to exploit makes it especially dangerous, as it can be weaponized by any attacker who can reach the vulnerable service over the network.

Action Items

- ☐ Apply the BeyondTrust security update for CVE-2026-1731 immediately and verify that all instances, including redundant or disaster recovery deployments, have been patched.
- ☐ Restrict network access to BeyondTrust management interfaces and agent communication ports, ensuring they are not exposed to the public internet and are accessible only from authorized management networks.
- ☐ Audit BeyondTrust session logs and administrative activity for signs of unauthorized access, unexpected command execution, or the creation of new accounts or access policies that were not initiated through approved change management processes.
- ☐ Conduct a broader review of all remote access and privileged access management tools in the environment to confirm that they are fully patched, properly segmented, and monitored with the same rigor applied to any other critical infrastructure component.

SECTION 4

Top 5 Social Engineering Attacks

AI rewrote the human attack surface. Voice phishing overtook email, and deepfakes moved from theoretical to multi-million-dollar reality

04

AI-Generated Phishing at Scale

4×

higher click-through vs. human-written lures

AI-generated phishing has moved from an emerging concern to a dominant force in the threat landscape over the course of Q2 2026. What began as a dramatic surge in late 2025, when AI-crafted phishing messages jumped from roughly 4% of observed attacks to over 50% during the holiday season, has now settled into a sustained baseline that shows no signs of retreating. The phishing messages being produced by generative AI tools are qualitatively different from their predecessors. They are grammatically polished, contextually aware, and capable of mimicking the tone and formatting of legitimate business communications with a degree of fidelity that renders many traditional detection cues obsolete. The spelling errors, awkward phrasing, and generic greetings that once served as reliable warning signs are increasingly absent.

The operational economics of AI-generated phishing have fundamentally shifted the advantage toward attackers. Crafting a convincing, personalized phishing message previously required manual research and writing effort that limited the volume of targeted campaigns. Generative AI eliminates that bottleneck entirely. Attackers can now scrape LinkedIn profiles, company websites, press releases, and social media activity, then feed that information into AI models that produce tailored lures at machine speed. The result is spear-phishing quality at mass-phishing volume. Research from multiple threat intelligence vendors has confirmed that AI-generated phishing lures achieve substantially higher click-through rates than their human-written counterparts, with some studies measuring engagement rates more than four times higher.

The challenge for defenders is that AI-generated phishing undermines the foundational assumption of most security awareness training programs: that users can be taught to spot phishing by looking for telltale signs of poor craftsmanship. When the craftsmanship is indistinguishable from legitimate communication, detection must shift from human pattern recognition to a combination of technical controls, process-based verification, and continuous behavioral training that goes beyond static red flags. Organizations that have not updated their phishing awareness programs to account for AI-generated content are training their employees to defend against a threat that has already evolved past the examples being used in the training.

Action Items

- ☐ Update security awareness training programs to include examples of AI-generated phishing that lack traditional red flags, emphasizing verification behaviors such as confirming requests through a separate communication channel rather than relying on visual inspection of the message itself.
- ☐ Deploy AI-capable email security solutions that analyze message content, metadata, and behavioral patterns to detect AI-generated phishing, rather than relying solely on signature-based or reputation-based filtering.
- ☐ Implement internal policies that require out-of-band verification for any email-initiated request involving financial transactions, credential changes, or sensitive data access, regardless of how legitimate the message appears.
- ☐ Conduct regular phishing simulations using AI-generated lures to establish a realistic baseline of organizational susceptibility and to identify employees or departments that may require additional targeted training.

Voice Phishing (Vishing) Overtaking Email

23%

of cloud-related compromises began with vishing

Voice phishing has undergone a dramatic escalation and has now displaced email as the leading social engineering vector in confirmed security incidents. Incident response data from Q2 2026 shows that email phishing dropped to just 6% of confirmed initial access methods, while vishing rose to 11% overall and reached 23% in cloud-related compromises. This inversion represents one of the most significant shifts in the social engineering landscape in recent memory. Attackers are no longer primarily relying on malicious links and attachments to gain footholds; they are picking up the phone, calling help desks and employees directly, and talking their way past identity verification processes in real time.

The effectiveness of modern vishing campaigns is being amplified by AI-powered voice cloning technology that has become dramatically more accessible and convincing over the past year. Attackers can now generate synthetic voice replicas of specific individuals using as little as a few seconds of publicly available audio sourced from corporate webinars, podcast appearances, LinkedIn videos, or conference recordings. These cloned voices are deployed in calls to help desks, finance teams, and IT support staff, where the attacker impersonates an executive or trusted colleague and requests credential resets, MFA bypasses, or urgent financial transfers. The quality of modern voice synthesis is sufficient to fool people who know the impersonated individual personally, and the calls often leave no malware trace, making forensic investigation significantly more difficult.

The rise of vishing as the dominant social engineering method exposes a gap in many organizations' security postures. Most security awareness programs and technical controls have been optimized for email-based threats: email filtering, link scanning, attachment sandboxing, and phishing simulation campaigns. Voice channels have received comparatively little investment in terms of both training and technical controls. Help desk procedures for identity verification are frequently weak, relying on easily obtainable information such as employee ID numbers, dates of birth, or manager names. Until organizations treat voice channels with the same rigor they apply to email, vishing will continue to be the path of least resistance for sophisticated attackers.

Action Items

- ☐ Overhaul help desk and IT support identity verification procedures to eliminate reliance on easily researched information, implementing callback verification to known-good numbers and requiring multi-factor confirmation before processing sensitive requests such as password resets or MFA changes.
- ☐ Extend security awareness training to include realistic vishing scenarios, ensuring that employees in high-risk roles such as help desk staff, executive assistants, and finance team members are trained to recognize voice-based social engineering tactics.
- ☐ Establish organizational code words or rotating verification phrases for authorizing high-value transactions or sensitive operations over the phone, particularly between executives and finance teams.
- ☐ Evaluate and deploy call authentication technologies where available, and work with telecommunications providers to implement STIR/SHAKEN or equivalent caller ID verification to reduce the effectiveness of spoofed calls.

Deepfake Video Impersonation in BEC

$25M+

lost in a single deepfake video-call incident

Deepfake video impersonation has transitioned from a theoretical concern to an active and growing threat vector in business email compromise campaigns throughout Q2 2026. Attackers are using commercially available and open-source AI tools to generate real-time video impersonations of executives, deploying them in video conference calls with finance teams, accounting departments, and other personnel who have authority to initiate wire transfers or approve large payments. The most prominent case to date involved a multinational company that lost over $25 million after a finance employee participated in a video call with what appeared to be the company's CFO and several other colleagues, all of whom were deepfake-generated impersonations.

The barrier to producing convincing deepfake video has dropped significantly. Attackers no longer need extensive footage of a target to produce a usable impersonation. A handful of publicly available photos and short video clips, readily sourced from corporate websites, social media profiles, and recorded conference appearances, can be sufficient to generate a real-time video avatar that passes casual visual inspection during a video call. The technology is particularly effective in the context of business video conferencing, where resolution is often limited, camera angles are fixed, and participants are conditioned to accept minor visual artifacts as normal connectivity issues. When combined with AI-cloned voice, the result is a multi-sensory impersonation that is extraordinarily difficult to distinguish from a legitimate interaction in real time.

The organizational exposure to deepfake video impersonation is concentrated in finance, accounting, and executive administration functions where high-value transactions are authorized through interpersonal communication. Traditional BEC defenses, which focus on email header analysis, domain reputation, and message content scanning, provide no protection against an attack that occurs entirely within a live video call. Defending against this threat requires a fundamental shift in how organizations verify the identity of participants in sensitive transactions, moving from visual and auditory recognition to process-based controls that do not depend on the perceived identity of the person making the request.

Action Items

- ☐ Implement mandatory dual-authorization protocols for all wire transfers and high-value financial transactions, requiring approval from at least two individuals through separate communication channels before funds are released.
- ☐ Establish out-of-band verification procedures for any financial instruction received during a video call, requiring that the request be confirmed through a separate, pre-established channel such as a direct phone call to a known number or an in-person confirmation.
- ☐ Reduce the public availability of executive video and audio content where practical, reviewing corporate websites, social media accounts, and recorded event appearances for material that could be used to generate deepfake impersonations.
- ☐ Train finance and accounting teams specifically on deepfake video risks, including demonstrations of current deepfake capabilities, and reinforce that visual or auditory recognition of a colleague on a video call is no longer a reliable basis for authorizing transactions.

Pretexting-Driven Business Email Compromise

50%+

of all social engineering incidents now pretexting-driven

Pretexting has overtaken traditional phishing as the most common social engineering method at the core of business email compromise, now accounting for more than 50% of all social engineering incidents. Unlike conventional phishing, which relies on a single deceptive message containing a malicious link or attachment, pretexting involves the construction of an elaborate false scenario through sustained, multi-turn conversation. Attackers research their targets extensively, studying the victim organization's vendor relationships, payment cycles, approval processes, and internal communication norms, and then engage finance professionals in back-and-forth email exchanges that build trust and credibility before ultimately requesting a fraudulent payment or sensitive data transfer.

The sophistication of modern pretexting campaigns has been dramatically amplified by generative AI. Attackers no longer need near-native fluency in the target's language to craft convincing CEO impersonations or vendor payment requests. AI tools produce grammatically flawless, contextually appropriate messages that adopt the tone, vocabulary, and formatting conventions of legitimate business correspondence. The grammar and tone anomalies that BEC training programs have historically used as detection cues are increasingly absent. Analysis from multiple email security vendors has found that a significant and growing proportion of BEC messages now contain AI-generated content, and the trend is accelerating.

The financial impact of pretexting-driven BEC remains among the highest of any social engineering category. Reported losses from BEC attacks run into the billions of dollars annually, with individual incidents frequently exceeding tens of thousands of dollars and sometimes reaching into the millions. The median amount extracted per incident continues to climb as attackers refine their targeting and timing. What makes pretexting particularly dangerous is that it often bypasses technical controls entirely. The messages contain no malware, no malicious links, and no attachments. They are simply well-crafted conversations between what appears to be a trusted business contact and an employee authorized to move money. The attack succeeds or fails entirely on the strength of the social engineering and the robustness of the organization's verification processes.

Action Items

- ☐ Implement mandatory verification procedures for all payment requests and changes to vendor banking details, requiring confirmation through a pre-established communication channel that is independent of the email thread in which the request was received.
- ☐ Deploy email security solutions with BEC-specific detection capabilities that analyze conversational patterns, sender behavior anomalies, and linguistic indicators of AI-generated content, rather than relying solely on attachment or link scanning.
- ☐ Conduct targeted training for finance, accounting, and procurement teams on pretexting tactics, emphasizing that the absence of malicious links or attachments does not mean a message is safe, and that sophisticated BEC attacks are designed to look exactly like normal business communication.
- ☐ Establish clear escalation procedures for any payment request that deviates from standard processes, involves urgency language, or comes from a new or recently changed email address, even if the request appears to originate from a known contact.

North Korean Deepfake IT Worker Infiltration

Insider

credentialed access obtained through the hiring process itself

North Korea has operationalized a social engineering campaign of a fundamentally different character than the other entries on this list. Rather than targeting organizations from the outside through phishing, vishing, or impersonation, the North Korean IT worker scheme embeds state-sponsored operatives directly into Western companies by having them apply for and secure legitimate remote employment using fabricated identities, falsified credentials, and AI-generated deepfake personas. These operatives pass through standard hiring processes, join corporate payrolls, and gain authorized access to internal systems, source code repositories, and sensitive business data as part of their normal job duties. The scheme generates illicit revenue for the North Korean regime while simultaneously providing espionage access to the environments in which these operatives are placed.

The scale and sophistication of this program have grown considerably. Operatives use AI-generated profile photos, deepfake video for remote interviews, and synthetic identity documents that are increasingly difficult to distinguish from legitimate credentials. They are often supported by facilitators in other countries who provide local addresses, proxy internet connections, and banking infrastructure to make the employment relationship appear geographically legitimate. Once placed, these operatives have been observed exfiltrating proprietary data, introducing backdoors into code repositories, and conducting reconnaissance of internal systems for potential future exploitation. The program has been flagged by multiple government agencies and threat intelligence organizations as a significant and growing concern.

What makes this threat uniquely challenging is that it exploits the trust inherent in the employer-employee relationship rather than targeting a technical vulnerability or manipulating an employee into a single compromising action. The operative is not an outsider attempting to break in; they are an authorized insider with legitimate credentials, system access, and a plausible reason to be interacting with sensitive resources. Traditional social engineering defenses, which focus on recognizing and resisting external manipulation, provide no protection against an adversary who has already been hired. Defending against this threat requires a shift in how organizations approach identity verification during hiring, access governance for remote employees, and ongoing monitoring of insider activity.

Action Items

- ☐ Strengthen identity verification procedures during the hiring process for remote positions, including live video interviews with identity document verification, background checks that cross-reference multiple independent sources, and verification of employment history through direct contact with previous employers rather than reliance on submitted references.
- ☐ Implement least-privilege access controls for all employees, particularly remote contractors and new hires, ensuring that access to source code repositories, internal tools, and sensitive data is granted incrementally based on demonstrated need rather than provided broadly at onboarding.
- ☐ Deploy behavioral monitoring and insider threat detection capabilities that can identify anomalous access patterns, unusual data transfers, and after-hours activity that may indicate unauthorized exfiltration or reconnaissance by a compromised or malicious insider.
- ☐ Coordinate with legal, HR, and security teams to establish clear policies and response procedures for suspected fraudulent employment, including preservation of forensic evidence and reporting to relevant law enforcement and intelligence agencies.

SECTION 5

"Fix-Type" Social Engineering

The social engineering techniques detailed in the preceding section represent the highest-volume and most impactful categories of Q2 2026. However, a parallel development deserves special attention: the rapid proliferation and evolution of "Fix-type" attacks, a family of browser-based social engineering techniques that coerce users into executing attacker-supplied commands under the guise of resolving an error, completing a verification, or fixing a system problem. These techniques share a common pattern in which the malicious action is disguised inside a workflow the user trusts and expects to be safe, and they are increasingly being adopted by both cybercriminal operations and nation-state groups. The three most significant variants observed in Q2 2026 are ClickFix, CrashFix, and ConsentFix.

ClickFix

ClickFix is the foundational technique from which the broader family of Fix-type attacks derives its name and operational pattern. First documented by Proofpoint researchers in 2024 through campaigns attributed to the threat actor TA571, ClickFix has grown from a niche social engineering trick into one of the most dominant initial access methods in the current threat landscape. A 2025 Microsoft report indicated that 47% of initial access incidents originated from the ClickFix technique, and that prevalence has held steady or increased through Q2 2026. Recorded Future's Insikt Group assessed in March 2026 that ClickFix will very likely remain the dominant initial access vector throughout the year.

The core mechanic of a ClickFix attack is deceptively simple. A user visits a compromised or attacker-controlled website and is presented with what appears to be a routine interaction: a CAPTCHA verification, a Cloudflare bot check, a browser error, or a system warning. Behind the visual lure, malicious code is silently copied to the user's clipboard. The page then instructs the user to "fix" the problem by opening the Windows Run dialog (Win+R), pasting the clipboard contents, and pressing Enter. The user believes they are completing a verification step or resolving a browser issue. In reality, they are executing a PowerShell command, mshta call, or other scripting payload that downloads and installs malware. Common payloads include Lumma Stealer, AsyncRAT, Xworm, Latrodectus, and NetSupport RAT.

What makes ClickFix particularly effective is that the user initiates the execution themselves. Because the victim is the one opening the Run dialog, pasting the command, and pressing Enter, the attack bypasses many automated security controls.

What makes ClickFix particularly effective is that the user initiates the execution themselves. Because the victim is the one opening the Run dialog, pasting the command, and pressing Enter, the attack bypasses many automated security controls. Email filters never see an attachment. Endpoint detection tools see a user-initiated action through a trusted system utility. Sandboxes have nothing to detonate. The attack surface is the user's trust in a familiar workflow, not a software vulnerability, and that trust is remarkably durable. The technique is also cross-platform: macOS variants instruct users to paste commands into Terminal, and a newer April 2026 variant documented by Jamf Threat Labs uses the applescript:// URL scheme to auto-open Script Editor with a pre-loaded script. ESET has confirmed that ClickFix affects all major operating systems.

ClickFix is delivered through a wide range of vectors. Phishing emails with links to ClickFix landing pages remain common, but attackers have expanded to compromised legitimate websites with injected JavaScript, malicious advertisements on streaming and pirated content sites, search engine optimization poisoning, malicious browser extensions, and even GitHub repositories with weaponized README files. The Center for Internet Security tracked ClickFix as responsible for over a third of all non-malware alerts against U.S. public sector targets in the first half of 2025, and the technique has been observed in campaigns attributed to nation-state actors including Kimsuky, MuddyWater, and APT28. MITRE has added T1204.004 (Malicious Copy and Paste) specifically to capture the ClickFix copy-paste pattern, recognizing it as a distinct technique worthy of its own classification.

The continued evolution of ClickFix is what makes it a persistent concern rather than a passing trend. New delivery variants emerge monthly, each shifting the execution method to evade the detection rules built for the previous iteration. In February 2026, Microsoft identified a variant that replaced the Run dialog with Windows Terminal, bypassing detections tuned to suspicious activity launched from the Run dialog. In the same month, a DNS-based variant was disclosed that uses nslookup commands to retrieve payloads through DNS responses rather than traditional web requests, blending malicious traffic into normal network activity. In March 2026, Atos researchers documented a variant using net use commands to map WebDAV shares and execute batch files, moving initial access away from commonly monitored scripting engines entirely. This variant successfully bypassed Microsoft Defender for Endpoint and was only detected through proactive threat hunting focused on the RunMRU registry key.

Action Items

- ☐ Restrict PowerShell and script execution for non-administrative users through Group Policy, Windows Defender Application Control, or equivalent endpoint management tooling, as this directly disrupts the most common ClickFix execution chain.
- ☐ Deploy detection rules targeting ClickFix execution patterns, including suspicious commands written to the Explorer RunMRU registry key, PowerShell or cmd.exe spawned from the Run dialog, and clipboard-to-execution sequences involving scripting hosts.
- ☐ Update security awareness training to include ClickFix-specific scenarios, emphasizing that legitimate websites and services will never ask users to paste commands into the Run dialog, Terminal, or File Explorer address bar. Train employees to recognize the pattern rather than memorize specific visual lures, as the lures change frequently.
- ☐ Implement network egress filtering and restrict outbound access to protocols, ports, and services that are not operationally required, with particular attention to legacy utilities such as finger.exe (TCP port 79) and unusual DNS query patterns that may indicate DNS-based payload staging.
- ☐ Incorporate ClickFix into phishing simulation programs, testing employees with realistic copy-paste scenarios rather than only traditional link-click or attachment-open simulations.

CrashFix

CrashFix is a ClickFix variant identified by Huntress in January 2026 and subsequently confirmed by Microsoft Defender Experts in February 2026. Microsoft described CrashFix as a "notable escalation in ClickFix tradecraft" that combines user disruption with social engineering to increase execution success while reducing reliance on traditional exploit techniques. Where standard ClickFix attacks rely on simulated errors and fake warnings to create urgency, CrashFix takes the manipulation a step further by deliberately crashing the user's browser to manufacture a genuine problem, then presenting a fraudulent fix that leads to malware execution.

Action Items

- ☐ Implement browser extension management policies that restrict installation to a curated allowlist of approved extensions, preventing users from installing unvetted extensions from public stores without administrative approval.
- ☐ Deploy endpoint detection rules for the specific CrashFix execution chain, including finger.exe being copied or renamed in the temporary directory, suspicious chrome.runtime connection flooding, and PowerShell or script execution initiated from the Run dialog immediately following a browser crash.
- ☐ Enable Attack Surface Reduction rules in Microsoft Defender or equivalent endpoint protection, particularly the rule that blocks executable files from running unless they meet prevalence, age, or trusted list criteria.
- ☐ Apply network egress filtering to block outbound finger protocol traffic (TCP port 79), which has no legitimate operational use in most enterprise environments and is being actively abused as a command-and-control channel.
- ☐ Conduct targeted awareness training that specifically addresses the CrashFix pattern: if a browser crashes and immediately presents a "fix" or "scan" prompt upon restart, the user should report the incident to IT rather than follow the instructions.

The attack begins with the installation of a malicious Chrome extension. The documented campaign used an extension called "NexShield

Advanced Web Guardian," which was hosted on the official Chrome Web Store and masqueraded as a clone of the legitimate uBlock Origin Lite ad blocker. The extension was downloaded at least 5,000 times before removal. Once installed, NexShield transmits a unique tracking ID to an attacker-controlled server and then waits. Sixty minutes after installation, the extension begins executing a denial-of-service attack against the victim's own browser by creating an infinite loop of chrome.runtime port connections that exhausts system resources and causes the browser to become unresponsive and crash. When the user force-quits and restarts their browser, a fake "CrashFix" security warning appears, claiming a security threat has been detected and instructing the user to run a scan by opening the Windows Run dialog and pasting a command from the clipboard.

The command copied to the clipboard abuses the legitimate Windows utility finger.exe, which is intended for retrieving user information from remote systems. The command copies finger.exe to the temporary directory, renames it, and uses it to contact attacker infrastructure and download the next-stage payload. The final payload is ModeloRAT, a Python-based remote access trojan that establishes command-and-control communication through periodic beacon requests and creates a Run registry entry to maintain persistence across reboots. A notable aspect of the campaign is that the malware only activates on domain-joined hosts, strongly suggesting that the threat actor behind CrashFix, tracked as KongTuke, is specifically targeting corporate environments rather than individual consumers.

The sophistication of CrashFix lies in the layered deception. Unlike standard ClickFix, which relies on a fake error that a cautious user might dismiss, CrashFix creates a real problem that demands a response. The browser genuinely crashes. The disruption is real. The frustration is real. When the "fix" appears, it arrives in a context where the user is already primed to take action to resolve an actual problem they are experiencing. This psychological escalation makes CrashFix meaningfully more effective than its predecessor at convincing users to follow through with the malicious instructions, and it represents a template that other threat actors are likely to replicate.

ConsentFix

ConsentFix represents the most conceptually dangerous evolution of the Fix-type attack family because it moves the entire attack into the browser context, eliminates the need for endpoint-level code execution, and bypasses phishing-resistant authentication controls including passkeys and FIDO2 hardware keys. Discovered and documented by Push Security in late 2025, ConsentFix blends ClickFix-style copy-paste social engineering with OAuth consent phishing to steal persistent access tokens for Microsoft accounts without ever requiring the victim to enter a password or complete an MFA challenge.

The attack begins when a user visits a compromised legitimate website, often discovered through a normal Google search rather than delivered via email, which immediately circumvents email-based anti-phishing controls. The site displays a fake Cloudflare verification prompt instructing the user to complete a series of steps to prove they are human. The user is directed to press a key combination that triggers a Microsoft OAuth authorization flow for Azure CLI, a first-party Microsoft application. A localhost URL containing an OAuth authorization code appears in the browser. The user is then instructed to copy this URL and paste it into a field on the attacker's page. Once pasted, the attacker captures the authorization code, exchanges it for a persistent OAuth token, and gains effective control of the victim's Microsoft account through the Azure CLI application. If the user already had an active Microsoft session, no login prompt or MFA challenge appears at any point in the attack.

ConsentFix challenges a set of assumptions that many security programs have treated as foundational.

Why ConsentFix is Hard to Stop

EDR is blind to it.

No malware, no script, no process creation. Endpoint tools have nothing to detect.

You can't block the app without breaking workflows.

Azure CLI is a first-party Microsoft application — standard OAuth restrictions don't apply, and blocking it risks disrupting legitimate dev and ops workflows.

Phishing-resistant MFA doesn't help.

This attack targets an authorization code inside an already-authenticated session, not a credential. No login means no protection.

The implications of ConsentFix are significant for several reasons. First, because the attack happens entirely within the browser and never touches the endpoint, endpoint detection and response tools have no execution event to detect. There is no malware, no script, no process creation. Second, because the targeted application is Azure CLI, a first-party Microsoft application, many of the controls that organizations use to restrict third-party OAuth integrations do not apply. Administrators cannot simply block Azure CLI without potentially disrupting legitimate developer and operations workflows. Third, because no login is required when an active session exists, phishing-resistant authentication methods that bind to the legitimate domain provide no protection. The attack does not phish a credential; it phishes an authorization code within an already-authenticated session. Since the initial research, the security community has identified several additional vulnerable Microsoft applications beyond Azure CLI that can be exploited through the same pattern.

ConsentFix challenges a set of assumptions that many security programs have treated as foundational. Organizations that invested heavily in passkeys and FIDO2 to eliminate credential phishing are not protected against an attack that harvests authorization codes from existing sessions. Organizations that locked down third-party OAuth applications are not protected against abuse of first-party applications. Organizations that focused detection on endpoint telemetry are not protected against an attack that never leaves the browser. ConsentFix forces a rethinking of identity security that extends beyond authentication and into the ongoing governance of OAuth grants, session management, and application-level access controls.

Action Items

- ☐ Audit OAuth application grants across Microsoft 365 and Azure AD environments, identifying and revoking access for any unused, unrecognized, or overly permissive application authorizations, with particular attention to Azure CLI grants from non-developer or non-administrator accounts.
- ☐ Implement conditional access policies that restrict Azure CLI and other high-privilege first-party application access to specific user groups, device compliance states, and network locations, reducing the population of users who can be exploited through this attack.
- ☐ Deploy monitoring for anomalous Azure CLI login events, particularly logins originating from users outside of developer and systems administrator roles, logins from unexpected geographic locations, and logins that occur without a preceding interactive authentication event.
- ☐ Train employees to recognize that legitimate verification and CAPTCHA flows will never ask them to copy and paste URLs containing authorization codes, localhost addresses, or lengthy parameter strings, and to report any such request to the security team immediately.
- ☐ Evaluate browser-level security controls, such as enterprise browser solutions or browser extensions that monitor and restrict clipboard operations, OAuth consent flows, and redirects to localhost URLs initiated by untrusted web pages.

SECTION 6

Why It Matters

The through line across every section is convergence.

The threats documented in this report are not isolated incidents. Taken together, they describe a threat landscape in which the assumptions that most security programs were built on are being systematically dismantled, and the speed at which attackers adapt continues to outpace the speed at which defenders respond.

The ransomware ecosystem has matured beyond recognition. Qilin, Cl0p, Akira, and DragonForce no longer operate like opportunistic criminal gangs. They run structured affiliate programs, maintain professional leak site infrastructure, and make deliberate decisions about which sectors to target based on an informed understanding of which victims face the steepest costs from disruption and data exposure. Cl0p has abandoned encryption entirely in favor of pure data-theft extortion, demonstrating that ransomware operators no longer need to lock systems to extract payments. DragonForce's cartel model decentralizes intrusion operations across a growing affiliate base while centralizing negotiation and leak infrastructure, making the group harder to disrupt and easier to scale. These are not groups that will age out of relevance. They are refining their business models in real time, and the organizations they target need to be doing the same with their defenses.

The vulnerability landscape this quarter reinforced a pattern that has been building for years: edge-facing infrastructure is the preferred entry point for the most capable attackers, and many organizations are failing to patch it fast enough.

The state-sponsored dimension adds a layer of risk that extends well beyond the victim organizations themselves. Salt Typhoon's pre-positioning within North American telecommunications infrastructure is not a conventional breach. It is a strategic investment in long-term access that can be activated for espionage, intelligence collection, or disruptive operations at a time of geopolitical advantage. The techniques Salt Typhoon uses, living off the land, exploiting edge-facing network devices, tunneling through legitimate protocols, are specifically designed to avoid detection for months or years. When the most sophisticated adversaries are embedding themselves inside the infrastructure that carries the communications of government agencies, defense contractors, and critical service providers, the consequences of delayed detection are national in scope.

The vulnerability landscape this quarter reinforced a pattern that has been building for years: edge-facing infrastructure is the preferred entry point for the most capable attackers, and many organizations are failing to patch it fast enough. CVE-2026-20127 in Cisco Catalyst SD-WAN was a maximum-severity authentication bypass that had been exploited since at least 2023 before it was publicly disclosed, and the attackers who used it demonstrated the operational discipline to downgrade firmware, establish persistence, and restore original software versions to cover their tracks. The cPanel/WHM bypass allowed a single ransomware group to compromise entire shared hosting environments at scale, affecting hundreds of downstream websites from each individual intrusion. The Fortinet and BeyondTrust vulnerabilities continued a multi-year trend of edge security appliances and privileged access tools being turned against the organizations they are designed to protect. YellowKey, the BitLocker bypass disclosed without a CVE or patch one day after Microsoft's May Patch Tuesday, demonstrated that even trusted native components of the operating system can become part of the attack surface in ways that existing detection models are not built to catch. Every one of these vulnerabilities targeted systems that sit at the boundary between an organization and the outside world, and every one of them was exploited before many organizations had completed their patching cycles.

Social engineering has undergone a qualitative shift that renders much of the conventional defense playbook inadequate. AI-generated phishing has eliminated the craftsmanship gap that once separated mass phishing from targeted spear-phishing. Voice phishing has overtaken email as the leading initial access vector, and the help desk procedures that most organizations rely on for identity verification were designed for a world where callers could not convincingly impersonate executives in real time. Deepfake video impersonation has moved from theoretical to operational, with multi-million-dollar losses already documented. Pretexting-driven BEC has doubled in frequency because attackers have learned that sustained, multi-turn conversations with finance professionals are more effective than single-message lures. North Korean deepfake IT worker infiltration has introduced a threat model in which the adversary does not break in from the outside but is hired through the front door and given legitimate credentials, system access, and a paycheck. Each of these developments individually would warrant organizational attention. Together, they represent a fundamental erosion of the trust signals that security programs have historically depended on.

The emergence and rapid evolution of Fix-type attacks illustrates how quickly a single social engineering innovation can reshape the threat landscape. ClickFix, which was first documented in 2024, accounted for 47% of initial access incidents by 2025 and has continued to grow through Q2 2026. It works because the victim executes the malicious command themselves, bypassing endpoint detection, email filtering, and sandboxing in a single stroke. CrashFix escalated the technique by manufacturing genuine browser crashes to make the social engineering more convincing. ConsentFix moved the entire attack into the browser, eliminated the need for endpoint code execution, and bypassed phishing-resistant authentication controls that many organizations deployed specifically to defend against the previous generation of credential theft. The Fix-type family demonstrates a broader truth about the current threat landscape: attackers are not simply finding new vulnerabilities in software. They are finding new vulnerabilities in the way people interact with software, and those vulnerabilities are proving harder to patch.

The traditional boundaries between these threat categories are dissolving, and the organizations that continue to treat them as separate problems will find themselves defending against yesterday's threat model while today's attackers have already moved on.

The through line across every section of this report is convergence. Ransomware groups are adopting nation-state tradecraft. Nation-state actors are pre-positioning in commercial infrastructure. Vulnerabilities in edge devices are being chained together and exploited for years before disclosure. Social engineering is being supercharged by AI and weaponized through workflows that users have been trained to trust. The traditional boundaries between these threat categories are dissolving, and the organizations that continue to treat them as separate problems will find themselves defending against yesterday's threat model while today's attackers have already moved on. The gap between attacker capability and defender response time is not just widening. It is being exploited as a deliberate part of the attack strategy. Closing that gap requires treating security as a continuous operational discipline, with investments concentrated where the evidence says the risk is highest: identity verification, zero-trust architecture, rapid patching of edge infrastructure, and resilience against the attacks that bypass technical controls entirely.

SECTION 7

07 How to Respond

A durable set of priorities for a landscape that no longer fits the old categories.

Q2 2026 marks the point where the threat landscape stopped fitting the categories we built our defenses around.

Ransomware groups now operate as disciplined enterprises shifting toward data-theft extortion. Nation-state actors like Salt Typhoon are pre-positioning inside the North American telecommunications infrastructure that carries everyone else's traffic. And the most consequential edge vulnerabilities of the quarter, from Cisco Catalyst SD-WAN to cPanel to YellowKey, were exploited faster than most organizations could patch them, with CVE-2026-20127 abused since at least 2023 before public disclosure.

Social engineering has been rewritten by AI in parallel: voice phishing has overtaken email as the leading initial access vector in confirmed incidents, deepfake video impersonation has already produced losses exceeding $25 million in a single incident, and Fix-type attacks including ClickFix, CrashFix, and ConsentFix are bypassing endpoint detection, email filtering, and even phishing-resistant authentication by turning trusted user workflows into the attack surface.

The common thread is convergence: criminal groups are adopting state-sponsored tradecraft, state actors are embedding in commercial infrastructure, and the lines between technical and human exploitation are dissolving. Organizations that treat these as separate problems are defending yesterday's threat model while attackers have already moved on. The only durable response is to treat security as a continuous operational discipline, with investment concentrated where the evidence points: identity verification, zero-trust architecture, accelerated patching of edge infrastructure, and resilience against the attacks that bypass technical controls entirely.

01 **Maintain an aggressive, risk-based patch management program that explicitly accounts for edge-facing infrastructure on an accelerated timeline.** The exploitation timelines documented across CVE-2026-20127, CVE-2026-41940, CVE-2026-21643, and CVE-2026-1731 show that attackers are weaponizing vulnerabilities within days of disclosure, and in the case of the Cisco Catalyst SD-WAN flaw attributed to UAT-8616, exploitation had been occurring since at least 2023 before public disclosure. Standard 30-day patching cycles are no longer sufficient for internet-facing systems. Extend this discipline to often-overlooked categories such as IoT devices, OT systems, network appliances, hypervisors, and embedded management interfaces.

02 **Enforce phishing-resistant MFA broadly, but recognize its limits.** The absence of MFA on Cisco ASA and AnyConnect VPNs was a primary enabler of Akira intrusions, and single-factor remote access has fueled initial access for nearly every ransomware group in this report. At the same time, ConsentFix and OAuth-based session theft demonstrate that authentication strength alone is not a complete defense. Pair MFA with conditional access policies, session monitoring, and regular OAuth grant audits to address attacks that bypass the login event entirely.

03 **Make identity security a foundational discipline, because attackers aren't breaking in, they're logging in.** Valid accounts and stolen credentials are the primary or significant initial access method for Akira, Cl0p, and DragonForce and fuel lateral movement for Salt Typhoon, while the North Korean IT worker scheme demonstrates that fully credentialed access can even be obtained through the hiring process itself. Defending against this requires treating identity as a continuous lifecycle: enforce least-privilege access, prune standing privileges through regular reviews, promptly disable accounts at offboarding, separate standing administrative rights from just-in-time elevation, protect credential stores through Credential Guard and LSASS hardening, eliminate weak credential formats like Cisco Type 7 passwords, and monitor authentication events for impossible travel, infostealer dumps, and behavioral anomalies.

04 **Modernize security awareness training to reflect the current social engineering landscape.** The shifts documented this quarter, including AI-generated phishing, vishing overtaking email, deepfake video in BEC, pretexting-driven multi-turn conversations, and the entire Fix-type family, all undermine the visual and grammatical cues that traditional training relied on. Training should emphasize verification behaviors and out-of-band confirmation rather than spotting bad spelling, and should be refreshed frequently to keep pace with technique evolution.

05 **Establish strong identity verification procedures at the help desk and during onboarding.** Help desk social engineering was the initial access point for the Marks & Spencer incident attributed to DragonForce, and fabricated identities are central to the North Korean IT worker scheme. Eliminate reliance on easily researched information (employee ID, manager name, date of birth) for sensitive requests like password and MFA resets, and require callback verification, manager attestation, or in-person confirmation for high-risk actions. Apply similar rigor to remote hiring through live video interviews, document verification, and independently sourced employment history checks.

06 **Require dual authorization and out-of-band verification for financial transactions and changes to vendor banking details.** Pretexting-driven BEC now accounts for the majority of social engineering incidents, deepfake video calls have already produced losses exceeding $25 million in documented cases, and Cl0p's direct outreach to executives illustrates how attackers target the people authorized to move money. No transaction or banking change should be approved on the strength of a single email, call, or video meeting, regardless of how convincing the participants appear.

07 **Build resilience around the assumption that breaches and data exposure will happen.** Backups need to be offline, immutable, and tested through actual restoration exercises, since Qilin, Akira, DragonForce, and Cl0p all systematically destroy shadow copies and recovery options. Beyond technical recovery, organizations need pre-drafted regulatory notification templates, customer communication plans, legal escalation paths, and tabletop-tested incident response playbooks that account for dual extortion, leak site publication, and executive-level extortion contact.

08 **Reduce the public footprint of executives, finance staff, and other high-value targets.** The same publicly available audio, video, and biographical information that enables convincing AI-generated phishing, voice cloning, and deepfake video impersonation can be curated and minimized. Review corporate websites, social media presences, recorded conference appearances, and employee LinkedIn profiles for material that could be repurposed by attackers, and apply data minimization principles where practical without disrupting legitimate business needs.

09 **Take inventory of and govern the use of browser extensions, OAuth applications, and other browser-resident integrations.** CrashFix abused a Chrome Web Store extension downloaded at least 5,000 times, and ConsentFix exploited first-party Microsoft OAuth applications to bypass passkeys and FIDO2 entirely. Restrict extension installation to an approved allowlist, audit OAuth grants regularly across Microsoft 365 and other SaaS platforms, and revoke unused or overly permissive authorizations. The browser is increasingly the operating system for enterprise work, and it needs to be governed accordingly.

10 **Apply zero-trust principles to internal network architecture, not just perimeter access.** The lateral movement techniques used across every ransomware group in this report, including PsExec over SMB, RDP, Kerberos ticket abuse, GPO-based payload distribution, and Salt Typhoon's SSH pivoting on network devices, all depend on flat, permissive internal networks. Segment networks meaningfully, isolate management planes for hypervisors and network devices, restrict administrative tooling to hardened jump hosts, and require explicit authorization for cross-segment access.

11 **Take vendor and supply chain risk seriously as a first-order concern.** Cl0p's serial zero-day campaigns against MFT platforms, Qilin's targeting of MSPs through ScreenConnect-themed lures, and the cascading impact of the cPanel exploitation on downstream hosting customers all demonstrate that the weakest link in your security posture may be a vendor you trust. Maintain an inventory of critical third parties, understand their patch management and breach notification commitments contractually, and develop contingency plans for vendor compromise scenarios.

12 **Strictly adhere to cybersecurity fundamentals and ensure all personnel undergo annual phishing and social engineering training.** Speak with your UltraViolet Cyber TAM Representative to schedule a live phishing engagement.

13 **Perform annual tech refresh reviews to gain a holistic understanding of your infrastructure.** Speak with your UltraViolet Cyber TAM Representative to schedule a Red Team or Purple Team engagement to gain insight into the vulnerabilities in your environment.

SECTION 8

Appendix

Full detail of the top 20 observed techniques per threat actor, Q2 2026 with field context.

Qilin

T1190 **Exploit Public-Facing Application** Qilin affiliates routinely gain initial access by exploiting vulnerabilities in internet-facing infrastructure, with a particular focus on Fortinet appliances (including CVE-2024-21762 and CVE-2024-55591), Citrix gateways, and exposed Remote Desktop Protocol services. This technique allows attackers to establish a foothold without any user interaction, making timely patching of perimeter devices a critical first line of defense.

T1566.002 **Phishing: Spearphishing Link** In addition to attachments, Qilin affiliates have used spearphishing emails containing malicious links to redirect victims to credential harvesting pages or malware download sites. Sophos MDR documented campaigns in which Qilin affiliates spear-phished managed service provider administrators through ScreenConnect-themed lures to gain access to downstream customer environments.

T1059.003 **Command and Scripting Interpreter: Windows Command Shell** Qilin affiliates use cmd.exe extensively to execute PsExec commands that distribute the encryptor binary across the targeted network. Observed command lines include PsExec invocations with password arguments and spread flags designed to propagate the ransomware to multiple network shares simultaneously from a single compromised host.

T1053.005 **Scheduled Task/Job: Scheduled Task** Qilin creates scheduled tasks for both persistence and execution, including tasks pushed via GPO and locally created tasks disguised as legitimate software. One documented example involved a scheduled task named "TVInstallRestore" configured to run at logon, mimicking TeamViewer to avoid suspicion from administrators reviewing task lists.

T1003.001 **OS Credential Dumping: LSASS Memory** Qilin employs an embedded Mimikatz module to dump credentials from the Local Security Authority Subsystem Service process memory. Harvested credentials are used to escalate privileges, move laterally across the network, and access additional systems and resources that would otherwise be inaccessible from the initial point of compromise.

T1548.002 **Abuse Elevation Control Mechanism: Bypass User Account Control** Qilin can bypass standard User Account Control protections by using stolen tokens to launch processes at an elevated security context. This allows the ransomware to execute privileged operations, such as disabling security tools and modifying system configurations, without triggering UAC consent prompts that might alert the user.

T1566.001 **Phishing: Spearphishing Attachment** Qilin has been delivered through malicious email attachments targeting specific organizations. These campaigns often precede the main intrusion by weeks or months, with the initial phishing serving as a reconnaissance or access-establishment phase that is later leveraged by the ransomware affiliate for deeper penetration.

T1059.001 **Command and Scripting Interpreter: PowerShell** PowerShell is a core component of Qilin's operational playbook, used for deploying ransomware payloads to VMware vCenter and ESXi servers, enumerating Active Directory environments, discovering domain users and groups, and executing post-compromise scripts. Qilin's reliance on PowerShell makes script block logging and constrained language mode enforcement valuable detection and prevention controls.

T1484.001 **Domain or Tenant Policy Modification: Group Policy Modification** Qilin has abused Group Policy Objects to push scheduled tasks and ransomware payloads to domain-joined endpoints at scale. This technique allows the attacker to leverage the organization's own Active Directory infrastructure as a distribution mechanism, ensuring that the ransomware reaches every machine in the domain without requiring individual access to each host.

T1547.001 **Boot or Logon Autostart Execution: Registry Run Keys / Startup Folder** Qilin establishes persistence by creating RunOnce autostart entries in the Windows registry that point to dropped copies of the ransomware binary in the Public folder. This ensures that if the system reboots before encryption is complete, the ransomware will re-execute automatically on the next logon cycle.

T1134 **Access Token Manipulation** Qilin uses its embedded Mimikatz module for token manipulation, allowing affiliates to impersonate other users and elevate their security context without needing to authenticate with stolen passwords directly. This technique enables the attacker to operate under the identity of domain administrators or service accounts with broad network access.

T1021.002 **Remote Services: SMB/Windows Admin Shares** Qilin embeds a copy of PsExec within its payload and places it in the %Temp% directory under a randomly generated filename. PsExec is then used to propagate the ransomware encryptor across the network via SMB administrative shares, enabling rapid lateral movement and simultaneous encryption of multiple hosts.

T1570 **Lateral Tool Transfer** Beyond PsExec-based propagation, Qilin affiliates transfer additional tools and secondary encryptor binaries across the environment using legitimate file transfer mechanisms. Cisco Talos documented the distribution of a second encryptor named encryptor_1.exe to multiple hosts using PsExec and SMB shares during the lateral movement phase.

T1490 **Inhibit System Recovery** Qilin systematically destroys recovery options by executing vssadmin.exe to delete all volume shadow copies, disabling High Availability and Distributed Resource Scheduler in vCenter clusters, and using WMIC to change the Volume Shadow Copy Service startup type to manual. These actions ensure that victims cannot restore encrypted data without paying the ransom or relying on offline backups.

T1685 **Disable or Modify Tools** Qilin actively terminates antivirus-related processes and services to prevent detection and interference during the encryption phase. The Qilin.B variant enhanced these capabilities, incorporating more aggressive techniques for identifying and neutralizing security tooling before the main payload executes.

T1070.004 **Indicator Removal: File Deletion** Qilin deletes itself from infected hosts after execution to complicate forensic analysis and reduce the likelihood that defenders will recover the ransomware binary for reverse engineering. This self-deletion behavior is common among modern ransomware families and is designed to slow the development of detection signatures and decryption tools.

T1486 **Data Encrypted for Impact** The core objective of Qilin's attack chain is domain-wide encryption. The ransomware employs AES-256 or ChaCha20 for file encryption and RSA-4096 or RSA-2048 to secure the generated encryption keys. Variants are capable of targeting Windows workstations and servers, Linux systems, and VMware ESXi hypervisors, enabling comprehensive encryption across heterogeneous environments.

T1489 **Service Stop** Before initiating encryption, Qilin terminates specific services on compromised hosts, including database services, backup agents, and email servers. This ensures that files held open by running applications are released and available for encryption, maximizing the impact of the attack and reducing the likelihood of partial encryption failures.

T1685.005 **Clear Windows Event Logs** After completing its operations, Qilin clears Windows Event Logs to remove forensic evidence of the attack. This anti-forensics technique hampers incident response by eliminating the log entries that defenders would use to reconstruct the attack timeline, identify compromised accounts, and determine the scope of the intrusion.

T1688 **Safe Mode Boot** Qilin can reboot targeted systems into Windows Safe Mode before executing encryption. Because most security tools and endpoint protection agents do not load in Safe Mode, this technique allows the ransomware to operate in an environment with significantly reduced defensive coverage, increasing the likelihood of successful encryption without detection or interference.

Cl0p

T1190 **Exploit Public-Facing Application** Exploitation of public-facing applications is Cl0p's signature initial access method and the technique most closely associated with the group's identity. Cl0p has conducted serial zero-day campaigns against enterprise managed file transfer platforms, including Accellion FTA (CVE-2021-27101 through CVE-2021-27104), Fortra GoAnywhere MFT (CVE-2023-0669), Progress MOVEit Transfer (CVE-2023-34362), and more recently CentreStack (CVE-2025-14611) and Oracle E-Business Suite (CVE-2025-61882). These campaigns compromise hundreds of organizations simultaneously through a single vulnerability, making this technique the cornerstone of Cl0p's mass-exploitation model.

T1566.001 **Phishing: Spearphishing Attachment** In earlier campaigns and for targets that are not reachable through MFT exploitation, Cl0p and TA505 employ large-scale spearphishing campaigns using macro-enabled documents. These emails deliver initial loaders such as Get2, which then install secondary payloads including SDBOT, FlawedAmmyy, and Cobalt Strike. While Cl0p has increasingly shifted toward zero-day exploitation, phishing remains a viable fallback initial access vector.

T1059.001 **Command and Scripting Interpreter: PowerShell** Cl0p uses PowerShell extensively for downloading additional payloads, executing commands on compromised systems, disabling security tools, and conducting reconnaissance. PowerShell is a preferred execution method because it is natively present on Windows systems, trusted by default, and provides direct access to .NET libraries and Windows Management Instrumentation, enabling a wide range of post-compromise actions without dropping additional binaries to disk.

T1505.003 **Server Software Component: Web Shell** Web shell deployment is a critical persistence and access mechanism in Cl0p's MFT exploitation campaigns. During the MOVEit campaign, Cl0p installed a custom web shell called LEMURLOOT on compromised Transfer servers, which was used to enumerate directory contents, exfiltrate files, and execute commands. During the Accellion campaign, the group deployed DEWMODE for similar purposes. These web shells provide persistent, covert access that survives reboots and allows the attacker to return to exfiltrate data over extended periods.

T1195 **Supply Chain Compromise** Cl0p's operational strategy relies on compromising widely used third-party software to reach large numbers of downstream victims from a single intrusion point. Rather than targeting individual organizations one at a time, the group identifies software used across entire industries and exploits trust relationships between vendors and customers. This supply chain approach enables Cl0p to scale extortion operations far beyond what traditional targeted ransomware campaigns can achieve.

T1078 **Valid Accounts** Cl0p leverages stolen or compromised credentials throughout its attack lifecycle. Credentials obtained through phishing, credential dumping, or purchased from initial access brokers are used to authenticate to victim environments, access file transfer platforms, and move laterally across networks. The use of valid accounts allows Cl0p operators to blend in with legitimate user activity and reduces the likelihood of triggering security alerts.

T1059.003 **Command and Scripting Interpreter: Windows Command Shell** Cl0p uses cmd.exe to execute commands directly on compromised systems, including launching the ransomware payload, deleting shadow copies, disabling recovery options, and terminating services. The Windows Command Shell serves as the primary command execution interface during the final stages of the attack chain, particularly during encryption and anti-recovery operations.

T1133 **External Remote Services** Cl0p targets vulnerable external remote services, including exposed RDP endpoints, VPN gateways, and enterprise collaboration platforms, to gain initial access to victim networks. Once authenticated through stolen credentials or exploitation, these services provide a direct pathway into the internal environment without requiring the victim to interact with a phishing email or malicious file.

T1486 **Data Encrypted for Impact** When Cl0p deploys its encryption payload, it uses AES for file encryption with RSA key wrapping and appends the ".clop" extension (or variants such as ".Cllp" and ".Cl0p") to encrypted files. The ransomware deliberately skips system-critical files to keep the operating system operational for ransom communication. However, in its more recent campaigns from 2023 onward, Cl0p has increasingly favored data-theft-only extortion without deploying encryption, making this technique less central to their current operations but still present in their capability set.

T1567 **Exfiltration Over Web Service** In campaigns targeting managed file transfer platforms, Cl0p takes advantage of the platforms' built-in file transfer capabilities to exfiltrate data. Because the exfiltration occurs through a service that is designed to transfer large volumes of files, the activity may not trigger data loss prevention alerts or network anomaly detections that would flag unusual outbound transfers through other channels.

T1490 **Inhibit System Recovery** Cl0p deletes volume shadow copies using vssadmin and disables Windows recovery options through bcdedit. These actions prevent victims from restoring encrypted files using built-in Windows recovery mechanisms, increasing the pressure to pay the ransom. This technique is a standard component of Cl0p's encryption workflow and is executed immediately before or concurrent with the encryption process.

T1003 **OS Credential Dumping** Cl0p operators use tools such as Mimikatz to extract credentials from LSASS memory and other credential stores on compromised systems. Harvested credentials, including domain administrator accounts and service account passwords, are used to escalate privileges and facilitate lateral movement across the victim's network. Credential dumping is typically performed early in the post-compromise phase to enable the broadest possible access before encryption begins.

T1027.002 **Obfuscated Files or Information: Software Packing** Cl0p employs software packing to obfuscate its ransomware binary and hinder static analysis by security researchers and automated detection systems. Packed samples require unpacking before their true functionality can be analyzed, which delays signature development and gives the ransomware additional time to operate before detection rules are updated.

T1041 **Exfiltration Over C2 Channel** Data exfiltration is the primary objective in Cl0p's modern operational model. The group exfiltrates sensitive data from compromised environments before encryption (or in lieu of encryption entirely) and threatens to publish it on their Tor-based leak site, "CL0P^_-LEAKS," if ransom demands are not met. In MFT exploitation campaigns, exfiltration is conducted through the compromised file transfer platform itself, leveraging its existing functionality to move data out of the environment.

T1562.001 **Impair Defenses: Disable or Modify Tools** Cl0p actively disables or uninstalls security products on compromised systems before deploying its encryption payload. This includes terminating antivirus processes, disabling endpoint detection and response agents, and using msiexec.exe to silently remove security software. By neutralizing defensive tools before encryption begins, Cl0p ensures that the ransomware can operate without interference.

T1489 **Service Stop** Before encryption, Cl0p terminates processes and services related to backup software, database engines, email servers, and security solutions. This ensures that files held open by running services are released and available for encryption, and that backup processes cannot create new recovery points during the attack. The list of targeted services is configurable and has expanded over successive Cl0p variants.

T1553.002 **Subvert Trust Controls: Code Signing** Cl0p has used digitally signed binaries to evade detection by security tools that trust signed executables. By obtaining or stealing valid code signing certificates, Cl0p can bypass application whitelisting controls, reduce the likelihood of triggering endpoint protection alerts, and increase the chances that the ransomware binary will execute successfully on protected systems.

T1140 **Deobfuscate/Decode Files or Information** Cl0p uses XOR-based decryption routines to decode strings at runtime, keeping sensitive configuration data such as targeted file extensions, excluded directories, and command-and-control parameters encrypted within the binary until execution. This technique prevents static analysis tools from extracting meaningful indicators of compromise from the ransomware sample without first reverse engineering the decryption logic.

T1614.001 **System Location Discovery: System Language Discovery** Cl0p checks the keyboard layout and active system font charset against Russian and Commonwealth of Independent States language identifiers before proceeding with encryption. If the system is configured with a CIS language, the malware deletes itself and terminates. This geofencing behavior is a common feature among ransomware families with suspected Russian-speaking operators and is designed to avoid attracting the attention of local law enforcement in the group's home region.

T1218.007 **System Binary Proxy Execution: Msiexec** Cl0p abuses the legitimate Windows Installer service (msiexec.exe) to disable security tools on compromised systems. Because msiexec is a trusted, signed Microsoft binary, its execution is less likely to trigger security alerts than direct process termination or registry manipulation. This living-off-the-land approach allows Cl0p to neutralize endpoint protection while maintaining a low detection profile during the pre-encryption phase.

Akira

T1078 **Valid Accounts** Compromised credentials are Akira's primary initial access method. The group acquires valid VPN credentials through purchase from initial access brokers, credential stuffing, or prior phishing campaigns, and uses them to authenticate directly to single-factor VPN gateways. Because the login appears legitimate, this access method generates minimal alerting and allows Akira to bypass perimeter defenses entirely without exploiting a software vulnerability.

T1059.001 **Command and Scripting Interpreter: PowerShell** Akira uses PowerShell scripts extensively for credential harvesting, privilege escalation, and post-compromise automation. PowerShell enables the group to interact with Active Directory, execute reconnaissance commands, download additional tools, and stage payloads without writing custom binaries to disk, reducing the forensic footprint and leveraging a trusted system utility for malicious purposes.

T1003.001 **OS Credential Dumping: LSASS Memory** Akira deploys Mimikatz to dump credentials from the Local Security Authority Subsystem Service process memory. Harvested domain credentials, including those of domain administrators and service accounts, are used to escalate privileges and authenticate to additional systems across the network. LSASS dumping is typically performed early in the intrusion to maximize the scope of accessible resources before defenders detect the compromise.

T1558 **Steal or Forge Kerberos Tickets** Akira has used scripts to dump Kerberos authentication credentials from compromised systems. Stolen Kerberos tickets enable the group to authenticate to network resources without needing plaintext passwords, and forged tickets (Golden or Silver Tickets generated via Mimikatz) can provide persistent, stealthy access to domain resources even after password resets, making this technique a powerful tool for maintaining long-term access.

T1018 **Remote System Discovery** Akira employs network scanning tools including Advanced IP Scanner and MASSCAN to identify live hosts, open ports, and accessible services within victim networks. This reconnaissance provides the group with a map of the environment that informs lateral movement decisions and ensures that the ransomware deployment reaches the maximum number of systems during the encryption phase.

T1133 **External Remote Services** Akira consistently targets externally accessible VPN appliances, particularly Cisco ASA and AnyConnect deployments lacking multi-factor authentication, as its preferred entry point. Once authenticated with stolen credentials, the VPN provides the attacker with direct network-level access to the victim's internal environment, effectively placing them inside the perimeter from the first moment of the intrusion.

T1059.003 **Command and Scripting Interpreter: Windows Command Shell** The Windows Command Shell is used throughout Akira's attack chain for executing ransomware payloads, deleting shadow copies, stopping services, and running lateral movement commands. Both the original Akira variant and the Megazord variant rely on cmd.exe for payload execution, making command-line logging and monitoring a critical detection opportunity.

T1555 **Credentials from Password Stores** In addition to LSASS dumping, Akira uses LaZagne to extract credentials stored in password managers, web browsers, Windows Credential Manager, and other local credential stores. This broad-spectrum credential harvesting approach ensures that the group captures not only Active Directory credentials but also application-specific passwords, API keys, and cached authentication tokens that may provide access to additional systems and services.

T1482 **Domain Trust Discovery** Akira uses built-in utilities such as Nltest and third-party tools like AdFind to enumerate Active Directory trust relationships within victim environments. Understanding domain trust configurations allows the group to identify pathways for lateral movement between domains, locate high-value targets such as domain controllers and file servers, and determine the full scope of the environment before deploying ransomware.

T1021.001 **Remote Services: Remote Desktop Protocol** After obtaining valid credentials, Akira uses Remote Desktop Protocol for interactive lateral movement across the victim network. RDP provides the attacker with full graphical access to compromised systems, enabling manual exploration, tool deployment, and configuration changes that would be more difficult to accomplish through command-line-only access.

T1021.002 **Remote Services: SMB/Windows Admin Shares** Akira uses PsExec to execute commands and deploy payloads on remote systems through SMB administrative shares. PsExec enables the group to propagate the ransomware encryptor across multiple hosts simultaneously, creating Windows services on target machines to execute the payload. This technique is a standard component of Akira's lateral movement and ransomware deployment workflow.

T1219 **Remote Access Tools** Akira deploys legitimate remote access utilities such as AnyDesk and PuTTY to maintain persistent, interactive access to compromised environments. These tools provide a fallback access mechanism that is independent of the initial VPN compromise and can survive credential resets or VPN account revocation. Because AnyDesk and PuTTY are legitimate software, their presence on systems is less likely to trigger security alerts.

T1486 **Data Encrypted for Impact** Akira's core ransomware payload encrypts files across victim environments as part of its double extortion operations. The original Akira variant targets Windows systems using C++ with ChaCha20 encryption, while the Linux variant targets VMware ESXi hypervisors. The Akira_v2 variant, written in Rust, and the Megazord variant provide additional cross-platform encryption capabilities. All variants share significant code overlaps with the Conti ransomware family.

T1489 **Service Stop** The Akira_v2 and Megazord variants terminate specific services on compromised hosts before initiating encryption. Targeted services typically include database engines, backup agents, email servers, and other applications that hold files open and would prevent encryption. Stopping these services ensures maximum file coverage during the encryption phase and disrupts business operations.

T1560.001 **Archive Collected Data: Archive via Utility** Before encryption, Akira uses utilities such as WinRAR to compress and archive stolen data in preparation for exfiltration. Archiving data into compressed packages reduces the volume of network traffic required for exfiltration and allows the group to stage complete data sets for transfer to external infrastructure in a more efficient and less detectable manner.

T1570 **Lateral Tool Transfer** Akira transfers tools, scripts, and ransomware binaries between compromised systems using PsExec, SMB shares, and remote access utilities. The ability to stage and distribute tooling across the environment allows the group to prepare multiple systems for simultaneous encryption and ensures that credential harvesting and reconnaissance tools are available wherever they are needed during the intrusion.

T1562.001 **Impair Defenses: Disable or Modify Tools** Akira actively disables or modifies security tools in victim environments to prevent detection and interference during the encryption phase. This includes terminating endpoint protection agents, disabling real-time monitoring, and modifying security configurations. By neutralizing defensive tools before deploying the ransomware payload, Akira ensures that encryption proceeds without automated intervention.

T1490 **Inhibit System Recovery** Akira's ransomware variants delete volume shadow copies to prevent victims from restoring encrypted files using built-in Windows recovery mechanisms. This anti-recovery step is executed automatically by the ransomware binary before or during the encryption process, ensuring that the only viable recovery options are offline backups or payment of the ransom demand.

T1531 **Account Access Removal** Akira deletes administrator accounts in victim networks prior to deploying encryption. By removing administrative access, the group hinders the victim's ability to respond to the attack in real time, locks out IT teams from critical systems, and creates additional pressure to pay the ransom. This technique is particularly disruptive in environments where administrative access recovery procedures are not well documented or tested.

T1567.002 **Exfiltration Over Web Service: Exfiltration to Cloud Storage** Akira exfiltrates victim data to cloud storage services using Rclone, a legitimate command-line program for managing files on cloud storage. Rclone supports a wide range of cloud providers and can be configured to transfer data over encrypted connections, making the exfiltration traffic difficult to distinguish from legitimate cloud synchronization activity. The exfiltrated data is then used as leverage in the group's double extortion model, with threats to publish files on their leak site if the ransom is not paid.

DragonForce

T1566 **Phishing** Phishing remains one of DragonForce's primary initial access vectors. Affiliates craft targeted spearphishing emails containing malicious attachments or links designed to deploy initial loaders or harvest credentials. These campaigns frequently use social engineering techniques tailored to specific individuals within the target organization, and in high-profile incidents such as the Marks & Spencer attack, the intrusion began with a social engineering phone call to an IT help desk, demonstrating that DragonForce affiliates blend email-based and voice-based phishing depending on the target.

T1078 **Valid Accounts** DragonForce affiliates frequently use compromised credentials obtained through phishing, credential stuffing, or purchase from initial access brokers to authenticate to victim environments. Valid account access is used to log into VPN gateways, RDP endpoints, and remote management platforms. Because the authentication appears legitimate, this technique generates minimal security alerting and allows the attacker to operate within the environment under the guise of an authorized user.

T1059.001 **Command and Scripting Interpreter: PowerShell** PowerShell is used extensively by DragonForce affiliates for downloading additional malware stages, executing reconnaissance commands, disabling security tools, and automating post-compromise activities. AttackIQ's emulation of DragonForce's attack chain specifically highlights the use of Invoke-WebRequest and DownloadData cmdlets within hidden PowerShell windows to retrieve secondary payloads, making PowerShell execution monitoring and constrained language mode enforcement critical detection and prevention controls.

T1105 **Ingress Tool Transfer** DragonForce replicates itself and transfers additional tools across compromised networks by downloading payloads from command-and-control infrastructure and copying files between internal systems. The ransomware copies itself to multiple directories and distributes tooling to additional hosts after gaining initial access. Both network-level and endpoint-level controls should be configured to detect the delivery of known malicious binaries and anomalous file transfers between internal systems.

T1021.001 **Remote Services: Remote Desktop Protocol** After obtaining valid credentials, DragonForce affiliates use RDP for interactive lateral movement across victim networks. RDP provides full graphical desktop access to compromised systems, enabling manual reconnaissance, tool deployment, data staging, and encryption preparation. The use of RDP is particularly common in engagements where Scattered Spider serves as the initial access broker, as the group heavily favors interactive access methods.

T1190 **Exploit Public-Facing Application** DragonForce affiliates actively exploit known vulnerabilities in public-facing systems, including unpatched web servers, VPN appliances, and remote management platforms. The group targets organizations that have failed to apply security updates for critical vulnerabilities, using internet-wide scanning to identify exposed systems at scale. This approach allows affiliates to gain initial access without relying on user interaction, making it particularly effective against organizations with immature patch management processes.

T1133 **External Remote Services** DragonForce targets externally accessible remote services, particularly RDP and VPN endpoints, as entry points into victim networks. In one documented incident response engagement, CSIS identified that the threat actor gained access by exploiting remote management software that had been installed by a previous hosting provider and was never removed when the organization switched providers. This highlights how legacy remote access tools and forgotten external services create unmonitored entry points that DragonForce affiliates actively seek out.

T1059.003 **Command and Scripting Interpreter: Windows Command Shell** The Windows Command Shell is used throughout DragonForce operations for executing ransomware payloads, running batch scripts, deleting shadow copies, stopping services, and performing system configuration changes. Command-line execution is the primary interface for the final encryption and impact phases of the attack, and monitoring for suspicious cmd.exe activity, particularly when spawned by unexpected parent processes, provides valuable detection opportunities.

T1106 **Native API** DragonForce's ransomware variants use Windows Native API calls, including CreateProcessA, to spawn new processes for payload execution. By interacting directly with the Windows API rather than relying exclusively on command-line utilities, the ransomware can perform operations with greater control and reduced visibility to security tools that primarily monitor command-line activity. API-level monitoring and behavioral detection are necessary to catch this execution method.

T1219 **Remote Access Tools** DragonForce affiliates deploy legitimate remote access software, particularly AnyDesk, to establish persistent access to compromised environments. AnyDesk provides a reliable backdoor that is independent of the initial access vector and can survive credential resets, VPN revocations, and other remediation actions taken by defenders. Because AnyDesk is a legitimate commercial tool, its presence on systems is less likely to trigger automated security alerts, making it an effective persistence mechanism.

T1003 **OS Credential Dumping** DragonForce affiliates use credential dumping tools, including Mimikatz and similar utilities, to extract credentials from LSASS memory, the Security Account Manager database, and other credential stores on compromised systems. Harvested credentials are used to escalate privileges to domain administrator level and authenticate to additional systems across the network, enabling the broad lateral movement required to stage a domain-wide encryption event.

T1082 **System Information Discovery** DragonForce gathers detailed system information from compromised hosts, including system configurations, user credentials, installed software, and hardware details. This reconnaissance data is logged to a file on the compromised machine and used to inform lateral movement decisions, identify high-value targets for encryption, and support further compromise of the network. The silent collection of system information is designed to avoid detection by security software.

T1486 **Data Encrypted for Impact** DragonForce's encryption capability is the core of its impact phase. The initial LockBit 3.0-derived variant uses AES-256 for file encryption with RSA-2048 key wrapping. The newer Conti V3-derived variant employs ChaCha8 for faster file encryption while retaining RSA for key protection. The ransomware encrypts data across endpoints, servers, network shares, NAS storage, databases, and virtual machines, including VMware ESXi environments. After encryption, ransom notes are dropped on disk and the desktop wallpaper is changed to display the ransom message.

T1489 **Service Stop** Before initiating encryption, DragonForce terminates processes and services that could interfere with the encryption process or hold files open. Targeted processes typically include database engines, backup agents, email servers, and security tools. Stopping these services ensures maximum file coverage during encryption and disrupts the victim's ability to maintain business operations during and after the attack.

T1070 **Indicator Removal** DragonForce actively removes indicators of compromise from compromised systems by deleting system logs, clearing event logs, and erasing forensic artifacts that could help security tools detect its presence. This cleanup activity extends to removing dropped files, clearing temporary directories, and using fileless execution techniques where possible. By eliminating forensic evidence, DragonForce hinders incident response and makes it more difficult for defenders to reconstruct the attack timeline and determine the full scope of the intrusion.

T1134 **Access Token Manipulation** DragonForce's ransomware requires administrator privileges to terminate processes and initiate encryption. Once the malware confirms it has elevated privileges, it attempts to escalate further to SYSTEM-level access by enumerating running processes, identifying one with SYSTEM-level privileges, and manipulating its access token. This token manipulation allows the ransomware to operate at the highest privilege level on the compromised host.

T1083 **File and Directory Discovery** DragonForce enumerates files and directories on compromised systems and accessible network shares to identify data for both exfiltration and encryption. The ransomware uses file extension matching to target documents, databases, virtual machine files, and other high-value data types while skipping system-critical files that would render the operating system inoperable. This selective targeting maximizes impact while preserving the victim's ability to communicate and pay the ransom.

T1490 **Inhibit System Recovery** DragonForce deletes volume shadow copies and disables system recovery options to prevent victims from restoring encrypted data without paying the ransom. This anti-recovery step is executed as part of the standard encryption workflow and is present in both the LockBit-derived and Conti-derived variants. By eliminating local recovery options, the group forces victims to rely on offline backups or ransom payment as their only paths to data restoration.

T1027 **Obfuscated Files or Information** DragonForce employs multiple obfuscation techniques to conceal its ransomware payload and hinder analysis. These include packing the executable, encrypting portions of the code, and obfuscating DLL names and function calls that are resolved dynamically at runtime. Notably, reverse engineering has revealed multiple text signatures written in Chinese within the code, adding complexity to attribution efforts. These obfuscation measures are designed to evade static analysis by security tools and delay signature development by researchers.

T1562 **Impair Defenses** DragonForce disables or removes antivirus software and other security tools on compromised systems to prevent detection and interference during the encryption phase. This includes terminating security-related processes, modifying security configurations, and using legitimate system utilities to uninstall endpoint protection products. The combination of defense impairment with log deletion and obfuscation creates a layered evasion strategy that is designed to keep the attacker undetected throughout the entire attack lifecycle.

Salt Typhoon

T1190 **Exploit Public-Facing Application** Exploitation of public-facing network infrastructure is Salt Typhoon's primary initial access method. The group targets known, unpatched vulnerabilities in edge devices including Cisco IOS and IOS XE (CVE-2023-20198, CVE-2023-20273, CVE-2018-0171), Ivanti Connect Secure and Policy Secure (CVE-2024-21887), and Palo Alto PAN-OS GlobalProtect (CVE-2024-3400). The FBI emphasized in February 2026 that basic configuration errors and unpatched systems provided the entry points. Salt Typhoon specifically targets Web Services Management Agent endpoints for exploitation and employs double URL encoding to obfuscate exploit traffic.

T1110.002 **Brute Force: Password Cracking** Salt Typhoon extracts password hashes and weakly encrypted credentials from dumped network device configurations and cracks them offline. Cisco devices using Type 7 password encryption are particularly vulnerable, as this encoding is trivially reversible. Cracked credentials are used to authenticate to additional devices and services across the network, expanding the group's access without triggering failed authentication alerts that brute-force attacks against live services would generate.

T1136 **Create Account** Salt Typhoon creates new Linux-level user accounts on compromised network devices by directly modifying /etc/shadow and /etc/passwd files. These accounts provide persistent access at the operating system level beneath the network management plane, making them invisible to administrators who only monitor device configurations through standard CLI or management interfaces. This technique represents a deep level of device compromise that most network monitoring tools are not designed to detect.

T1562.004 **Impair Defenses: Disable or Modify System Firewall** Salt Typhoon modifies access control lists and loopback interface configurations on compromised network devices to create pathways for their operations. By altering ACLs, the group can open access from their controlled infrastructure to management interfaces that would normally be restricted, enable traffic flows that support data exfiltration, and create exceptions that allow their tools and tunnels to operate without being blocked by the device's own security controls.

T1602.002 **Data from Configuration Repository: Network Device Configuration Dump** Once Salt Typhoon compromises a network device, one of their first actions is to dump the device's running and startup configurations. Configuration files contain critical intelligence including network topology information, routing tables, access control lists, SNMP community strings, and encrypted or hashed credentials for local and remote authentication. This technique serves a dual purpose: it provides credentials for further compromise and gives the group a detailed map of the victim's network architecture.

T1098.004 **Account Manipulation: SSH Authorized Keys** Salt Typhoon establishes persistence on compromised network devices by adding SSH authorized keys under root or other user accounts at the Linux level. This technique allows the group to maintain remote access to the device even if passwords are changed, TACACS+ or RADIUS authentication is modified, or the device is rebooted. Because SSH key-based authentication is a legitimate and expected mechanism on network infrastructure, the added keys can be difficult to detect without systematic configuration auditing.

T1021.004 **Remote Services: SSH** Salt Typhoon uses SSH for lateral movement between compromised network devices. The group modifies loopback addresses on compromised switches and uses them as the source of SSH connections to additional devices, allowing traffic to bypass access control lists that would otherwise block connections from unauthorized source addresses. By enabling SSH on non-standard ports and sourcing connections from trusted internal interfaces, Salt Typhoon creates lateral movement pathways that are extremely difficult to distinguish from legitimate administrative traffic.

T1572 **Protocol Tunneling** Salt Typhoon modifies device configurations to create Generic Routing Encapsulation (GRE) tunnels and IPsec tunnels on compromised network infrastructure. These tunnels provide encrypted, covert command-and-control channels that traverse the victim's network using legitimate routing protocols. Because GRE and IPsec are standard networking protocols used in normal operations, tunnel traffic blends with legitimate network management activity and is unlikely to trigger anomaly-based detection without deep packet inspection and baseline comparison.

T1040 **Network Sniffing** Network sniffing is one of Salt Typhoon's most strategically significant techniques. The group uses a variety of tools and techniques to capture packet data between network interfaces on compromised routers and switches, including capturing TACACS+ and RADIUS authentication traffic to harvest credentials in transit. By intercepting traffic at the network infrastructure level, Salt Typhoon gains access to call records, subscriber metadata, communications content, and authentication credentials that pass through the compromised devices, providing intelligence collection capabilities far beyond what endpoint compromise alone would yield.

T1560 **Archive Collected Data** Salt Typhoon uses JumbledPath and other tools to archive collected packet captures and configuration data on compromised devices before exfiltration. Archiving reduces the volume of data that must be transferred and allows the group to stage complete intelligence packages for efficient extraction. The archived data includes captured network traffic, authentication credentials, configuration files, and other intelligence gathered during the operational period.

T1590.004 **Gather Victim Network Information: Network Topology** Salt Typhoon systematically maps victim network topology using information extracted from compromised device configurations. Routing tables, OSPF and BGP neighbor relationships, VLAN assignments, and interface configurations provide a comprehensive picture of upstream and downstream network segments, peering relationships, and interconnections with other providers. This intelligence is critical for identifying additional targets, planning lateral movement, and understanding the scope of communications that pass through the compromised infrastructure.

T1078 **Valid Accounts** Salt Typhoon leverages compromised credentials extensively for lateral movement, using valid accounts obtained through configuration dumps, password cracking, and network sniffing to authenticate to additional network devices and systems. Because the group uses legitimate credentials rather than exploits for much of their lateral movement, their activity appears as authorized administrative access in device logs, making it extremely difficult to distinguish from normal operations without behavioral analysis and baseline comparison.

T1587.001 **Develop Capabilities: Malware** Salt Typhoon has developed custom tooling tailored to their operational requirements, most notably JumbledPath, a custom ELF binary designed for Cisco network devices. JumbledPath enables packet capture on compromised devices while simultaneously obfuscating the source of the capture request by routing traffic through multiple hops. This custom capability demonstrates the group's investment in developing purpose-built tools that are unlikely to be detected by signature-based security products.

T1048.003 **Exfiltration Over Alternative Protocol: Exfiltration Over Unencrypted Non-C2 Protocol** Salt Typhoon exfiltrates configuration files and other collected data from compromised network devices using FTP and TFTP, standard network management protocols that are commonly used for legitimate device configuration backup and software updates. By using these protocols rather than custom exfiltration channels, the group's data theft activity blends with normal network operations and is less likely to trigger data loss prevention or network anomaly alerts.

T1070.002 **Indicator Removal: Clear Linux or Mac System Logs** Salt Typhoon meticulously clears forensic evidence from compromised network devices by deleting .bash_history, auth.log, lastlog, wtmp, and btmp files. This log clearing removes evidence of the group's commands, authentication events, login timestamps, and other artifacts that incident responders would use to reconstruct the intrusion timeline. Combined with the group's use of living-off-the-land techniques that generate minimal malware-based indicators, log clearing makes detection and attribution significantly more difficult.

T1071.001 **Application Layer Protocol: Web Protocols** For command-and-control communications, Salt Typhoon uses application layer protocols including HTTPS to blend C2 traffic with normal web activity. The group has been observed using Cobalt Strike beacons and custom implants that communicate over encrypted web channels, and maintaining dual-channel C2 approaches that combine dedicated attacker infrastructure with legitimate services such as AnonFiles, File.io, GitHub, and Gmail. This diversified C2 architecture provides redundancy and makes disruption of the group's communications difficult.

T1588.002 **Obtain Capabilities: Tool** In addition to custom-developed malware, Salt Typhoon uses publicly available tools and exploit code to complement their operational toolkit. This includes publicly available exploit scripts for targeting known CVEs in network infrastructure, as well as open-source utilities for reconnaissance, credential harvesting, and network manipulation. The use of public tools alongside custom capabilities creates a mixed tooling profile that complicates attribution and makes it difficult for defenders to distinguish Salt Typhoon activity from other threat actors.

T1059 **Command and Scripting Interpreter** Salt Typhoon's post-compromise operations rely heavily on living-off-the-land techniques using native command interpreters and scripting capabilities available on compromised systems. On network devices, this includes using Cisco IOS CLI commands, Linux shell access on the underlying operating system, and SNMP SET requests for configuration changes. On Windows hosts within victim environments, the group uses PowerShell and WMIC for lateral movement, reconnaissance, and tool execution. The reliance on native utilities rather than custom malware minimizes the group's detection footprint.

T1053 **Scheduled Task/Job** Salt Typhoon uses scheduled tasks to establish persistence and ensure recurring execution of their tools and implants on compromised systems. On Windows hosts within victim environments, the group has been observed using schtasks to create scheduled tasks that execute PowerShell or cmd.exe commands, providing a reliable mechanism for maintaining access and re-establishing C2 communications if sessions are interrupted. On network devices, equivalent persistence mechanisms ensure that the group's access survives reboots and configuration reloads.

T1665 **Hide Infrastructure** Salt Typhoon uses JumbledPath and multi-hop routing techniques to obscure the origin of their operations and hide their command-and-control infrastructure. By routing traffic through multiple compromised devices and using legitimate network protocols for tunneling, the group makes it extremely difficult for defenders to trace activity back to its true source. This infrastructure obfuscation is a hallmark of Salt Typhoon's operational security and reflects the group's emphasis on long-term, undetected access over rapid exploitation and impact.

ultraviolet

Close the gap between tested and actually detected.

Speak with your UltraViolet Cyber TAM to schedule a live phishing engagement, or a Red or Purple Team assessment, and gain insight into the vulnerabilities in your environment.

Schedule My Assessment

About UltraViolet Cyber

UltraViolet Cyber is the only security operations partner that unifies red, blue, and purple team capabilities into one integrated team — finding what's vulnerable, stopping active threats, and validating that your defenses hold under real pressure. Built by former U.S. intelligence community operators with 30+ years of experience, we serve 400+ Global 2000 enterprises and federal agencies. Our practitioner-led and AI-accelerated closed-loop operations turn offensive findings into defensive weapons immediately — so gaps close in real time and defenses improve before attackers can exploit what testing uncovers. Offense informs defense. Defense sharpens offense. Security that gets smarter — and stronger — with every iteration.

UltraViolet is at the forefront of AI security — purpose-built to test, validate, and govern the AI systems organizations are deploying today. From adversarial testing of models, pipelines, and agents to pre-production threat modeling and board-ready governance frameworks, UltraViolet brings the same offensive and defensive rigor to AI that it applies across the enterprise. UltraViolet is ranked #19 on the Top 250 MSSP List and is headquartered in McLean, Virginia.

443.351.7630 / info@uvcyber.com | UltraViolet Cyber | @uv_cyber