June | 2026

Arete

Crimeware Report

Cyber Threat Trends and Highlights | Q1 2026

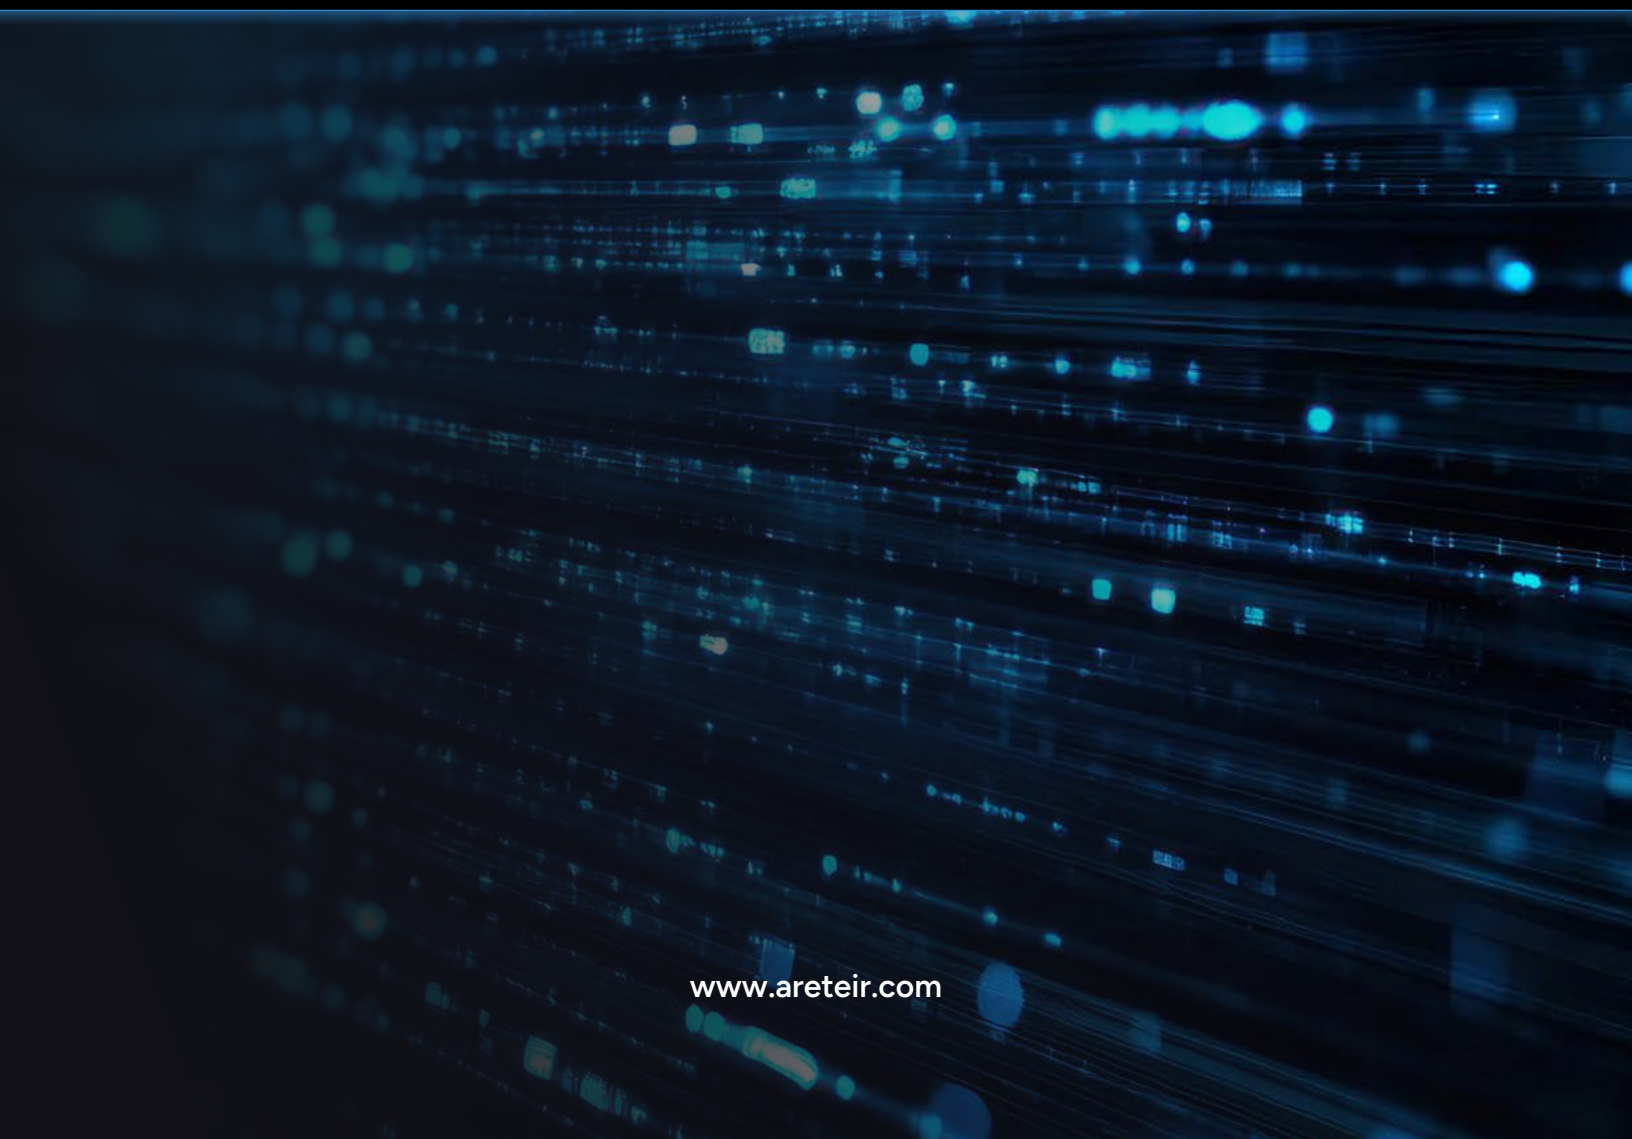

www.areteir.com

Table of Contents

Scan QR code to access the 2025 Annual Crimeware Report

Overview

Arete helps organizations stay ahead of cyber threats with standardized, intelligence-led detection, response, resolution, and resilience capabilities. Leveraging data and intelligence collected during ransomware and extortion incident response engagements, this report highlights notable trends and shifts in the threat landscape during the first quarter (Q1) of 2026.

Many of the trends observed at the end of 2025 continued into Q1 of 2026. Akira remained the top ransomware threat throughout Q1, followed by Qilin, as has been the case since July 2025. Trends in ransom payments, impacted industries, and initial access also remained relatively consistent with those observed at the end of 2025.

However, there were still notable shifts in the threat landscape. Threat actors continued to adapt and refine social engineering tactics in an ongoing shift toward identity-driven compromise techniques over traditional credential theft. Additionally, the role of AI in the cyber threat landscape is expanding, and threat actors notably leveraged AI tools to accelerate analysis of exfiltrated data in Q1.

The Akira and Qilin ransomware groups accounted for almost a third of all ransomware and extortion incidents in Q1 2026. Four ransomware groups—Akira, Qilin, PLAY, and INC Ransom—have remained among the top five threat groups since the second half of 2025.

The median ransom demand in Q1 was lower than in 2025, but the median payment increased, which can be partially attributed to the larger volume of Akira attacks observed.

The sectors most impacted by ransomware and extortion events in Q1 also remained relatively consistent with the end of 2025, and threat actors continued to target access vectors and vulnerabilities over specific business sectors.

Vulnerability exploits, compromised credentials, and social engineering attacks were the top initial access vectors observed in Q1. Threat groups continue to have success with social engineering tactics, including the reemergence of Microsoft Teams–based social engineering campaigns, device code phishing, and ClickFix-style lures.

Statistics and Trends from Arete's Incident Response Engagements

Q1 2025

33
Total Named Threat Actors

9
Total Unnamed Threat Actors

Q1 2026

38
Total Named Threat Actors

10
Total Unnamed Threat Actors

The threat landscape in Q1 2026 was primarily dominated by the Akira and Qilin ransomware groups, whose combined activity accounted for almost a third of all ransomware and extortion engagements that Arete responded to.

Akira and Qilin Remain Dominant Threats

As was the case in the second half of 2025, the threat landscape in Q1 2026 was primarily dominated by the Akira and Qilin ransomware groups, whose combined activity accounted for almost a third of all ransomware and extortion engagements that Arete responded to. However, the percentage of attacks attributed to Akira and Qilin did start to decline, as these groups were collectively responsible for over half of engagements in both Q3 and Q4 of 2025.

Akira's activity, in particular, began to normalize this year, with monthly activity returning to levels similar to those observed prior to the group's July 2025 surge. This decline was somewhat expected after Arete observed an unprecedented, historically high volume of activity from Akira in Q3 last year, largely due to a SonicWall VPN access control vulnerability. Although Akira was still responsible for over a third of all engagements in Q4 2025, the downward trend indicated that the group's late summer surge was unsustainable. Regardless, from July 2025 through the end of Q1 2026, Akira accounted for a larger share of activity each month than any other threat group observed by Arete.

TOP THREAT GROUPS OBSERVED BY QUARTER

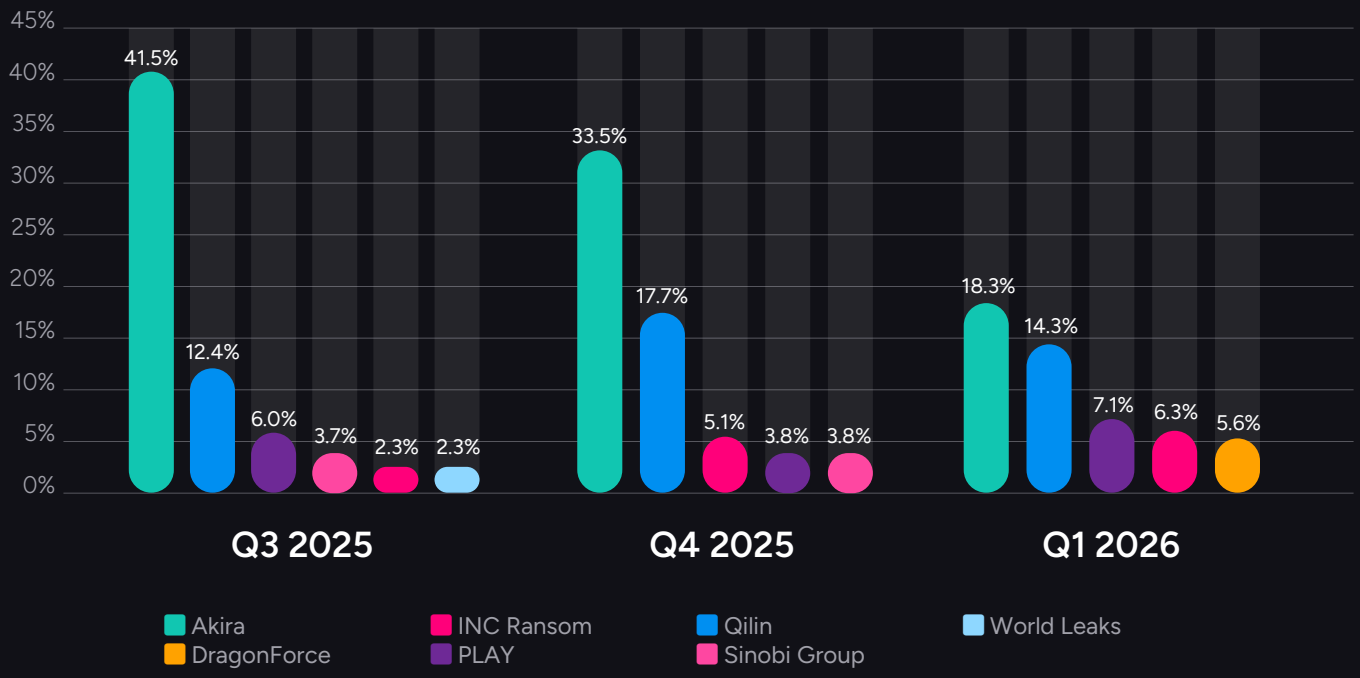

TOP THREAT GROUPS BY MONTH

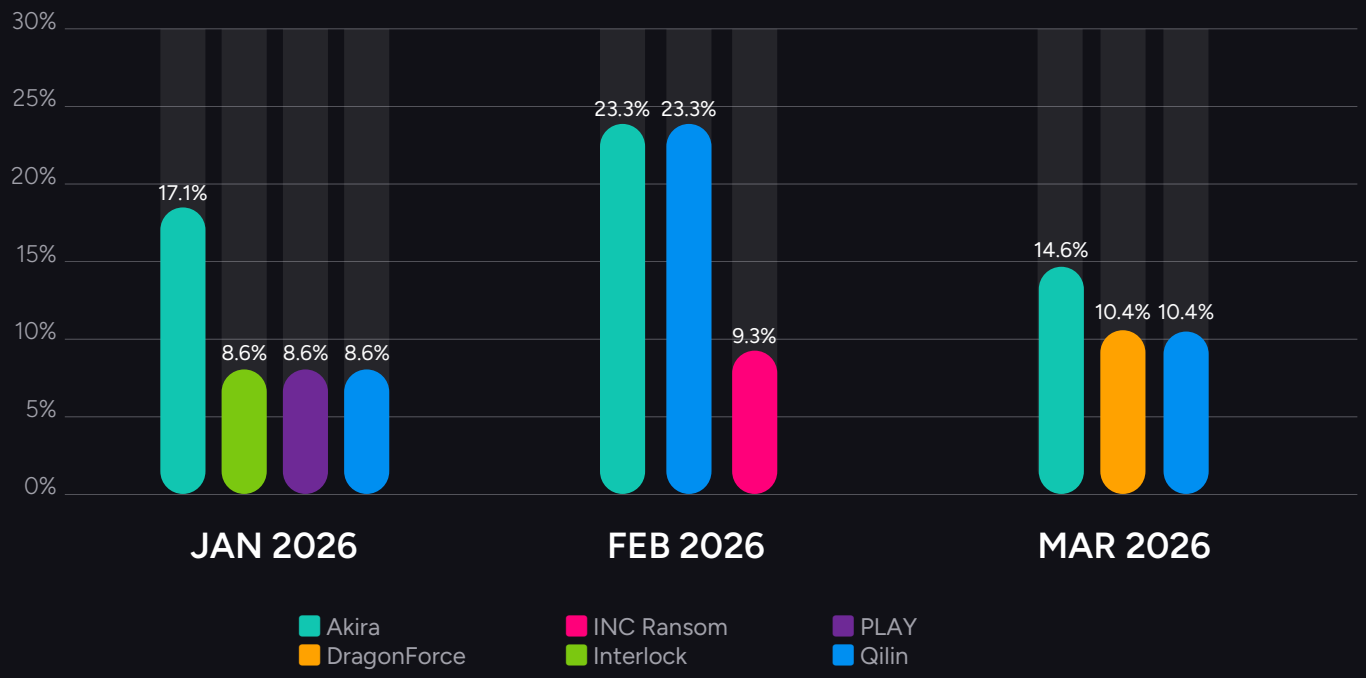

Qilin's Ransomware-as-a-Service (RaaS) operation remained the second most active ransomware group in Q1, behind Akira. Qilin had a notable uptick in activity in February, during which Arete observed Qilin actively targeting WatchGuard Firebox devices to gain initial access, especially devices vulnerable to CVE-2025-14733, a critical vulnerability in WatchGuard Fireware OS that allows a remote, unauthenticated threat actor to execute arbitrary code.

The rest of the Q1 threat landscape also remained relatively predictable. Despite the ongoing decline in Akira engagements, the majority of attacks continue to be conducted by the same established threat groups. Four groups—Akira, Qilin, PLAY, and INC Ransom—have been among the top five ransomware threats each quarter since the second half of 2025. Barring any unforeseen internal or external disruption to these groups, it is likely this trend will continue throughout the second quarter of the year.

New and Emerging Threat Groups in Q1

Although established threat groups conducted a majority of ransomware and extortion activity in Q1, Arete also observed several newcomers to the threat landscape.

NightSpire

NightSpire is believed to be a closed, internally managed ransomware group. Although the group launched its dark web data leak site (DLS) in March 2025, NightSpire has been substantially more active this year, posting over 100 victim disclosures in Q1 alone. NightSpire is believed to be a rebrand of the Rbfs ransomware group, based on overlapping victim disclosures, shared infrastructure, and the abrupt disappearance of Rbfs as NightSpire became active.

BravoX

BravoX is an emerging RaaS operation that first appeared publicly in January 2026, following the publication of a TOR-based DLS and an announcement post on the RAMP underground forum. The group appears to be in an early credibility-building phase, with limited but deliberate activity rather than high-volume campaigns, and it posted just nine victims on its DLS in Q1 2026. Despite the low victim count, BravoX's deployment of a standalone leak platform signals long-term intent rather than short-term monetization, but it remains to be seen whether the group will be successful in attracting affiliates or increasing its operational tempo this year.

Schrodinger Cat

Schrodinger Cat is a ransomware sub-group first observed in February 2024 and associated with the GlobeImposter ransomware family. In Q1 2026, the group appeared to be formalizing its brand, launching a DLS with victims dating back to July 2025. The group primarily targets enterprise environments, encrypting data and appending the .schrodingercat extension to affected files.

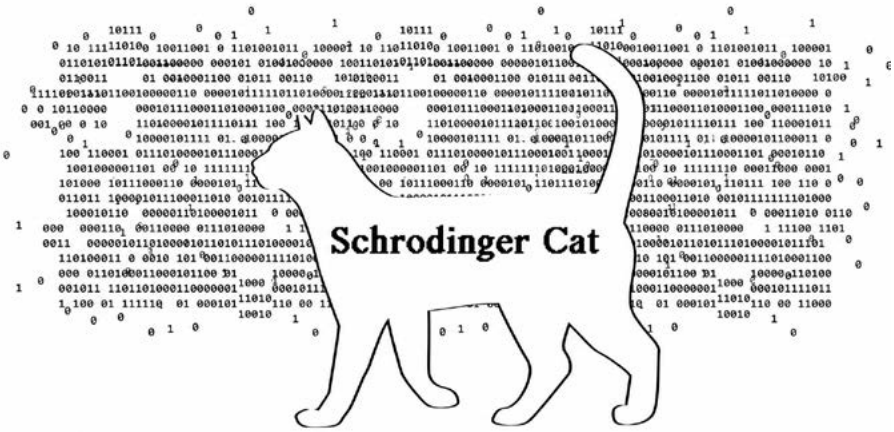

Figure 1 | Image from Schrodinger Cat's DLS (Source: Arete)

Threat Actor Spotlight:

DragonForce Cartel

The DragonForce ransomware group has had a notable uptick in activity so far in 2026. Although the group has operated since at least August 2023, in March of 2025, it announced that it was rebranding as the DragonForce Ransomware Cartel, reportedly transitioning from a centralized RaaS structure to a decentralized partnership model. Under this new model, affiliates have access to DragonForce tools and infrastructure, with the option to either operate under their own brands or deploy DragonForce ransomware.

In addition to the new cartel model, DragonForce has been linked to other threat groups and affiliates known to operate within multiple RaaS organizations. Although DragonForce incidents made up less than 1% of engagements Arete responded to in 2025, the group gained media attention last year when members of the Scattered Spider collective reportedly used DragonForce ransomware in attacks against high-profile UK retail companies Marks & Spencer and Harrods.

Much like its decentralized organizational structure, DragonForce's ransomware also reflects an evolving ecosystem influenced by the adoption of leaked source code, code reuse, and affiliate-driven customization.

The group initially used a variant of the leaked LockBit 3.0 builder in 2023 before launching its RaaS in 2024. This new DragonForce ransomware is likely based on Conti's 2022 leaked source code, which multiple threat actors have leveraged to accelerate development. Malware analysis by Arete's Threat Research Team identified multiple similarities between the leaked Conti code and DragonForce's ransomware, including an identical ChaCha-like ARX encryption routine, identical command-line arguments, and a mutex value in the 32-bit version. For technical details, see Appendix: Code Comparison, Lineage, and Evolution of Conti, DragonForce, and Devman Ransomware.

Ultimately, DragonForce's emergence as a self-described cartel underscores how threat groups continually evolve their strategies to maximize profits in an ever-changing threat landscape. Unlike traditional RaaS models, in which core developers establish dark web infrastructure and recruit affiliates from the cybercrime underground, DragonForce's structure gives affiliates greater flexibility, enabling the group to remain a persistent threat in 2026.

DragonForce has been linked to other threat groups and affiliates known to operate within multiple RaaS organizations.

Devman: From DragonForce Affiliate to Short-lived RaaS

First appearing at the beginning of 2025, a threat actor self-identified as "Devman" initially operated as an affiliate of multiple established RaaS groups, including Qilin and DragonForce. In April 2025, Devman began independently deploying a lightly modified version of the DragonForce encryptor, though early versions inadvertently encrypted the ransom notes as well, leaving victims without a means to communicate with the threat actor. By September 2025, Devman formally launched a RaaS operation, recruiting affiliates on the dark web and replacing the initial "Devmans Place" DLS with "Devman 2.0."

Despite launching an independent ransomware brand, Devman's ransomware continued to exhibit overlaps with DragonForce's ransomware and, by extension, Conti's leaked source code. The 32-bit Devman ransomware is functionally and structurally identical to DragonForce, with shared components including command-line parameters, mutex values, process and service kill lists, file and folder exclusions, encryption workflow, and decryption logic.

In January 2026, Devman announced a proposed capability, "Devman 3.0," described as a channel intended to facilitate contact between the group and potential insider sources. The concept appeared to be designed to lower barriers for individuals who may be willing to sell data or access but lack the technical knowledge or resources required to engage directly via the dark web or encrypted criminal platforms.

However, shortly after this announcement, Devman activity declined abruptly. Public communications ceased across social media and dark web infrastructure, and the most recent victim entry on the group's DLS was posted in early February 2026. Since then, no confirmed activity has been observed across known Devman-associated personas or platforms. Despite this apparent shutdown, indicators of possible follow-on activity have emerged. As early as April 2026, security researchers identified ransomware samples deploying a ransom note consistent with Devman-associated branding. This raises the possibility that a successor variant, Devman 2.1, may be in the final stages of development. Any potential resurgence appears notably more subdued than prior campaigns. There has been no observed revival of Devman's public branding, affiliate recruitment efforts, or social media presence. Whether this reflects increased operational caution, activity by a different actor reusing existing tooling, or continuation of previously established infrastructure remains unclear.

Devman's journey from affiliate to standalone RaaS brand to facilitator for other cybercriminals demonstrates how the cyber extortion landscape continuously evolves, necessitating an intelligence-led approach to threat actor attribution in any given cyber incident.

Trends in Ransom Demands and Payments

As cyber incidents increasingly intersect with financial crime, Arete uses threat intelligence, malware reverse engineering, and advanced blockchain analysis to drive defensible decisions about ransom payments, which are reflected in our data.

Trends in ransom payments remained largely consistent throughout Q1 2026. The median ransom demand for Q1 increased compared to the same timeframe in 2025, but was lower than the median demand for all of 2025. The median payment, on the other hand, increased to $250,000 in the first quarter of this year, compared to $100,000 in Q1 2025 and $152,750 for all of 2025.

This increase in median payment can be partly attributed to the high volume of Akira engagements observed by Arete. Although Akira's activity continued to decline during Q1 compared to Q3 and Q4 of 2025, the group was still responsible for almost 20% of all observed activity.

Akira's median demand and payment were higher than the collective median, at $700,000 and $300,000, respectively, for the quarter. The median payment made to Akira was also higher than that of any of the other top five threat groups observed in Q1 (Qilin, PLAY, INC Ransom, and DragonForce).

The percentage of time a ransom is paid decreased slightly to 31% in Q1 2026, compared to the first quarter and the overall percentage for 2025. While Akira saw success in collecting ransom payments during Q1, several other top threat groups had little success, despite their attack volume. No ransom payments were made to the DragonForce ransomware group in Arete engagements during Q1, and PLAY collected a ransom in only 11% of engagements, largely due to aggressive pressure tactics and an unwillingness to negotiate, coupled with victims' ability to recover without paying for decryption.

Overall, Arete has observed that the percentage of time a ransom is paid remains generally consistent quarter to quarter, and this will likely continue in the near term.

Akira collected an unusually high percentage of ransom payments in Q1, with victims opting to make payments to Akira almost 48% of the time. Of the payments made to Akira, over half of victims only paid for data suppression and not the decryptor, likely contributing to the group's higher payment percentage for the quarter.

PERCENTAGE OF TIME A RANSOM IS PAID

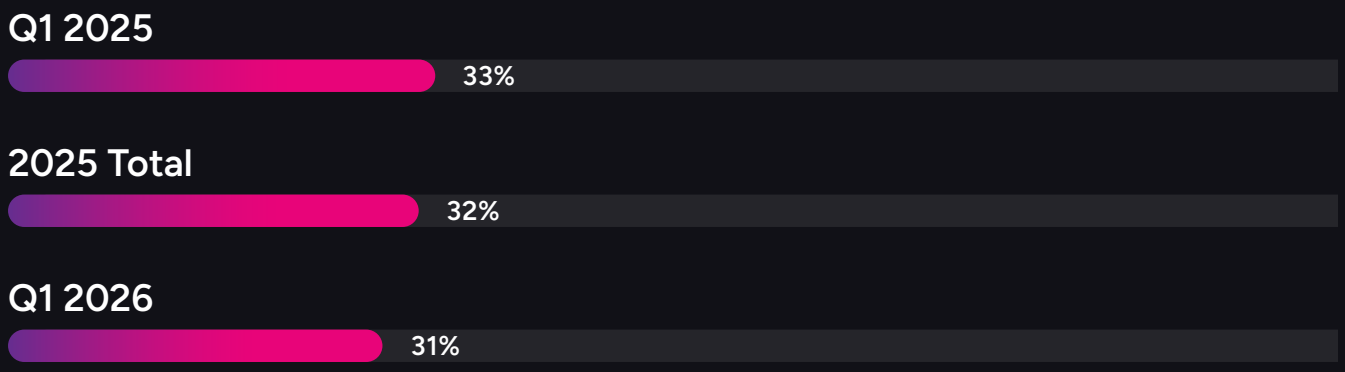

MEDIAN RANSOM DEMAND

MEDIAN RANSOM PAYMENT

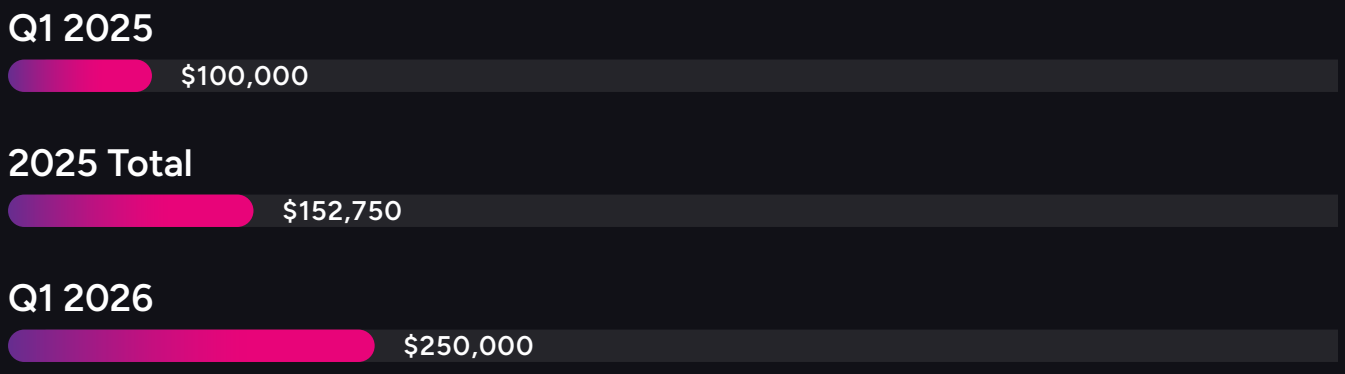

Sector Impacts and Threat Actor Targeting

The top sectors impacted by ransomware and extortion events in Q1 remained relatively consistent with those at the end of 2025, with Professional, Scientific, & Technical Services remaining the most impacted industry, accounting for a little over a fifth of all engagements, followed by Manufacturing, Wholesale Trade, and Healthcare & Social Assistance.

Consistency in this data over time indicates that the majority of active threat groups remain opportunistic, targeting specific initial access vectors rather than specific sectors or organization types. Groups like Luna Moth, which continued to largely target law firms in Q1 2026, remain the exception and not the norm. Instead, high-activity groups like Akira and Qilin often target specific vulnerabilities in perimeter devices rather than pursuing victims in any particular sector.

Conversely, threat actors' increasing focus on high-impact supply chain attacks against Software-as-a-Service (SaaS) platforms, open-source ecosystems, MSPs, and cloud integrations is likely to have a greater effect on individual sectors, given the disproportionate way these incidents act as force multipliers that affect downstream organizations.

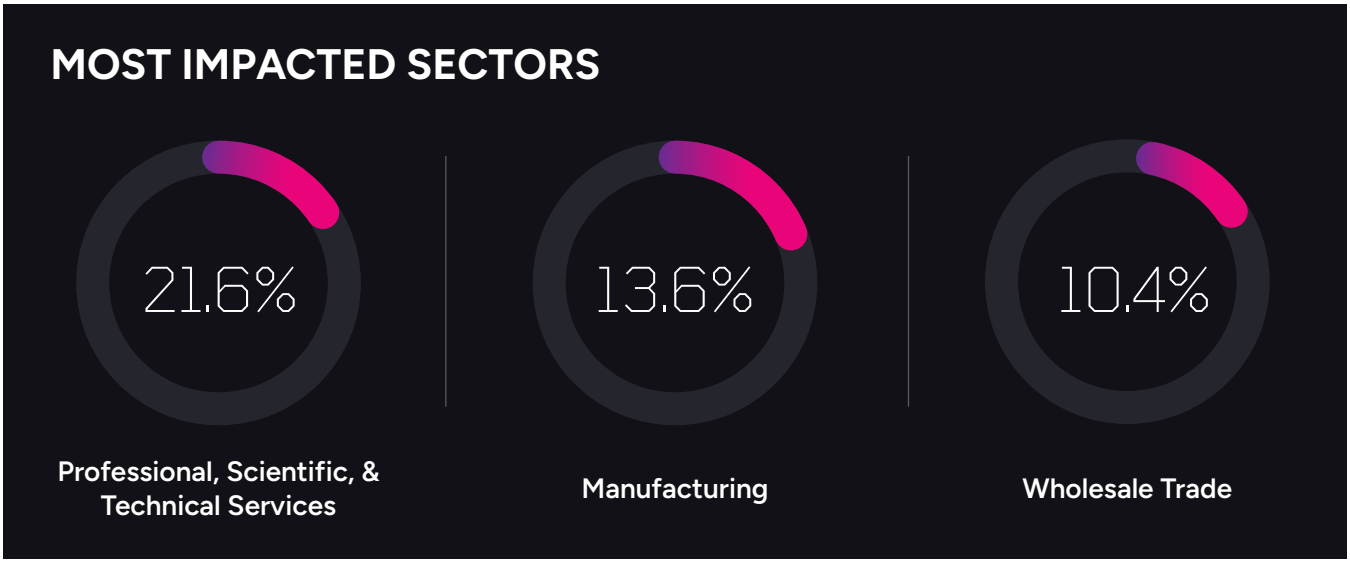

NAICS SECTOR NAME*	PERCENTAGE OF ENGAGEMENTS
Professional, Scientific, & Technical Services	21.6%
Manufacturing	13.6%
Wholesale Trade	10.4%
Healthcare & Social Assistance	8.8%
Retail Trade	8.0%
Construction	7.2%
Other Services (except Public Administration)	5.6%
Educational Services	4.8%
Finance and Insurance	3.2%
Public Administration	2.4%
Transportation & Warehousing	2.4%
Uncategorized	1.6%
Management of Companies & Enterprises	1.6%
Arts, Entertainment, & Recreation	1.6%
Administrative & Support & Waste Management & Remediation Services	1.6%
Information	1.6%
Real Estate & Rental & Leasing	1.6%
Agriculture, Forestry, Fishing & Hunting	0.8%
Mining, Quarrying, & Oil & Gas Extraction	0.8%
Accommodation & Food Services	0.8%

Figure 2 | Most impacted NAICS sectors in 2025 (Source: Arete)

** The North American Industry Classification System (NAICS) is the standard used by Federal agencies to classify US business organizations. The Cybersecurity and Infrastructure Security Agency (CISA) has their own separate classifications of critical infrastructure sectors.*

Initial Access Trends and Vulnerabilities Exploited

As was the case at the end of 2025, vulnerability exploits, compromised credentials, and social engineering attacks remain the top initial access vectors so far in 2026. Firewalls, in particular, continue to be prime targets for threat actors due to their strategic position at the network edge. Because these systems are directly exposed to the internet and are often integrated with authentication mechanisms, their compromise can provide a highly effective entry point into internal environments.

Additionally, as threat actors shift away from traditional credential phishing toward identity-driven compromise techniques, social engineering continues to evolve. What began as simple impersonation has transformed into coordinated campaigns leveraging OAuth phishing, device-code authentication abuse, token theft, and session hijacking to gain persistent access to victim environments. These attacks are increasingly scalable and automated, often enhanced with AI-generated content, enabling adversaries to convincingly impersonate internal users and bypass traditional security controls.

Firewalls continue to be prime targets for threat actors due to their strategic position at the network edge.

As threat actors shift away from traditional credential phishing toward identity-driven compromise techniques, social engineering continues to evolve.

Top Vulnerabilities Exploited in Q1

Fortinet Devices Targeted

Fortinet appliances, particularly FortiGate Next-Generation Firewalls, remained a prominent initial access vector in the first quarter of 2026. Threat actors actively targeted these devices using a combination of critical vulnerability exploitation and credential-based attacks. Several large-scale campaigns during this period compromised hundreds of Fortinet devices worldwide in a short timeframe. Notably, these operations also increasingly incorporated AI-assisted tooling to enhance reconnaissance, automate vulnerability identification, and streamline exploitation, reflecting a broader shift towards more scalable and efficient attack methodologies.

The primary driver behind these compromises was the exploitation of high-severity vulnerabilities, particularly those enabling authentication bypass and pre-authentication access. Key vulnerabilities observed during this period included CVE-2025-59718 and CVE-2025-59719, which allowed SAML authentication bypass, CVE-2026-24858, affecting FortiCloud SSO, and CVE-2026-35616, enabling pre-authentication API access bypass. In addition to targeting newly disclosed flaws, threat actors also demonstrated opportunistic behavior by exploiting older, unpatched vulnerabilities such as CVE-2019-7192, CVE-2023-27532, and CVE-2024-40711. This pattern indicates widespread internet scanning and targeting of exposed devices regardless of patch levels. In several cases, initial exploitation was followed by the establishment of secondary access channels, including Remote Desktop Protocol (RDP), to facilitate deeper network penetration.

Alongside vulnerability exploitation, credential-based attacks against Fortinet devices remained a consistent and effective secondary access vector. Threat actors targeted VPN portals and administrative interfaces through Botnet-driven brute-force campaigns, attempting common usernames like “admin”. They also leveraged credential reuse and datasets harvested from infostealer malware, significantly improving the success rate of authentication-based attacks. Ransomware groups, including Qilin, Safepay, DragonForce, and Embargo, were observed combining these techniques during Q1, using vulnerability exploitation to gain an initial foothold and brute-force attacks to expand access through VPN services.

Because Fortinet devices are often integrated with enterprise authentication systems such as Active Directory and Lightweight Directory Access Protocol (LDAP), their compromise can provide threat actors with extensive visibility and control across the network, significantly increasing the risk of lateral movement and data exfiltration. Analysis of incident response data further revealed notable variation in attacker dwell time, ranging from immediate post-compromise activity to delays of up to two months before further exploitation. Some campaigns prioritize quick monetization, such as ransomware deployment shortly after gaining access, while others focus on maintaining persistent, low-profile access for future use.

WatchGuard Fireware OS Vulnerabilities

During Q1 2026, multiple critical vulnerabilities were identified in WatchGuard Technologies' Firebox appliances running Fireware OS, significantly increasing the risk to organizations relying on these devices for perimeter security. As internet-facing firewall and VPN systems, these appliances represent high-value targets, and the discovery of severe flaws capable of enabling unauthenticated remote code execution (RCE) further amplified their exposure.

The most severe, CVE-2025-14733, is a critical vulnerability that allows an unauthenticated attacker to execute arbitrary code and establish persistent control over the affected system. Other vulnerabilities, including CVE-2025-9242, further contribute to the risk by enabling remote code execution, particularly in older versions such as Fireware OS 12.8.2. Collectively, these flaws create a critical security gap in network edge infrastructure, where successful exploitation can directly expose internal enterprise environments. As mentioned earlier in this report, the Qilin ransomware group was observed exploiting these WatchGuard vulnerabilities in Q1 and incorporating them into intrusion workflows to gain initial access and facilitate ransomware deployment in enterprise environments.

BeyondTrust Vulnerability

BeyondTrust solutions became a significant focal point for threat actors in Q1 2026 following the widespread exploitation of CVE-2026-1731, a critical vulnerability affecting BeyondTrust Remote Support and Privileged Remote Access products. As these solutions are commonly deployed to manage and secure privileged access, their exposure to the internet makes them highly attractive targets. This pre-authentication RCE vulnerability enabled attackers to compromise systems without valid credentials, making it a highly effective initial access vector, particularly in ransomware-driven intrusions.

The observed attack chain typically begins with scanning for exposed BeyondTrust instances, followed by exploitation using crafted requests to achieve unauthenticated remote code execution. Once access is established, attackers deploy reverse shells using tools such as bash, Python, and netcat to maintain control, and in several cases, leverage additional tools like VShell for remote access and SparkRAT for post-exploitation control and persistence. The impact of successful exploitation can be severe, often resulting in the compromise of remote access infrastructure, unauthorized control over internal systems, credential theft, and widespread network intrusion.

Arete incident response engagements and forensic investigations confirmed active exploitation of this vulnerability during the first quarter, and a threat actor cluster associated with Medusa ransomware operations was linked to this activity.

Axios NPM Compromise

In late March, a state-sponsored threat actor compromised the popular npm library “axios” by introducing two malicious dependency packages, which included a post-installation script that installed a cross-platform remote access tool (RAT). The compromise was facilitated through stolen maintainer credentials, which allowed attackers to publish malicious versions of the package directly to the npm registry. This approach bypassed traditional code review and Continuous Integration and Continuous Delivery/Deployment (CI/CD) security controls, demonstrating how trust in widely adopted libraries can be exploited to gain initial access.

Despite partial execution failures in some cases, the potential impact of this supply chain attack is significant. By targeting developer environments before application deployment, threat actors gained unauthorized remote access, introduced malicious code into production pipelines, and enabled widespread compromise across applications that depended on the affected package. This incident underscores the systemic risk associated with compromised dependencies in modern software development, as well as a growing shift in threat actor focus toward open-source ecosystems and trusted development dependencies. By targeting widely used libraries such as axios, threat actors were able to position themselves upstream in the software development lifecycle, enabling potential downstream impact at scale.

Notable Social Engineering Tactics Observed in Q1

Microsoft Teams Social Engineering Attacks Reemerge

A social engineering tactic appears to have re-emerged in 2026 in which threat actors impersonate IT staff on Microsoft Teams to gain initial access. The now-defunct Black Basta ransomware group was first observed using this tactic in late 2024, when a threat actor would flood a user's inbox with thousands of non-malicious emails and then message the user on Microsoft Teams posing as the victim organization's IT Help Desk and attempt to convince them to install a remote desktop support tool, giving the threat actor a means of initial access. After Black Basta shut down its brand in early 2025, other groups like Cactus were observed using the same tactic, but the campaign seemed to die out after Q1 of 2025.

However, in Q1 2026, multiple threat groups were once again observed using similar Microsoft Teams–based social engineering attacks, with the technique becoming more refined, targeted, and operationally mature. Threat actors are now increasingly abusing cross-tenant collaboration features to impersonate internal support functions and convince users to grant access through legitimate workflows. Attackers still commonly impersonate internal roles such

as IT support or help desk staff, using external tenant messaging or guest access in Teams to appear legitimate. Campaigns often begin with phishing emails or email bombing, followed by Teams-based interactions that guide victims through OAuth consent flows or remote access sessions, including the abuse of tools like Quick Assist.

So far this year, the Akira, INC Ransom, Payouts King, and Chaos ransomware groups have been observed using this tactic for initial access and intrusion operations. These campaigns highlight the shift toward identity-first attacks and the exploitation of trust in collaboration platforms and authentication systems. By leveraging legitimate services and mimicking normal user behavior, attackers are able to evade traditional defenses and maintain persistent access. The recent reemergence of these tactics underscores the need for strong identity security controls, continuous monitoring of authentication activity, and enhanced user awareness to defend against increasingly sophisticated social engineering threats.

The Evolution of ClickFix Tradecraft

ClickFix social engineering campaigns continued to evolve in Q1 2026, driven by increasingly sophisticated user deception techniques, expanded cross-platform targeting across Windows and macOS, and adoption by ransomware groups including Akira, Qilin, and Interlock. By leveraging legitimate system tools and relying on user interaction rather than software exploits, ClickFix has proven highly effective at bypassing traditional security controls.

Since 2025, ClickFix attacks have increasingly used realistic, seemingly urgent prompts, such as fake CAPTCHA pages, counterfeit Blue Screen of Death errors, browser alerts, and system troubleshooting messages. Notable developments in Q1 included the “CrashFix” campaign, which uses a fake adblocker that crashes the user's web browser and then tricks users into executing a Python-based RAT, and “InstallFix”, which leverages fake install websites promoted as sponsored results on search engines, which then trick users into installing malicious malware and infostealers.

In addition to these new deception techniques, ClickFix campaigns expanded significantly across platforms in Q1, with attackers extending operations into macOS environments. Campaigns grouped into distinct clusters based on infrastructure reuse, lure themes, and targeting strategies have also increased. These include impersonation campaigns targeting platforms such as QuickBooks and Booking.com, large-scale themed lures, dual-platform delivery models, and macOS-focused attack chains.

Defensive measures are beginning to evolve in response to ClickFix threats, with platform-level controls emerging as an important line of defense. A notable example is Apple’s introduction of a security feature in macOS Tahoe 26.4 that warns users when attempting to paste suspicious

commands into Terminal, effectively blocking execution. While such advancements represent meaningful progress in mitigating user-driven attacks, ClickFix remains a high-impact threat due to its reliance on human interaction, use of legitimate tools, and adaptability across environments, reflecting a broader shift toward deception-driven, low-code intrusions.

AI-Enhanced Cyberattacks Reshaping the Threat Landscape

In 2026, threat actors have increasingly leveraged AI in their operations, enabling rapid analysis of massive volumes of stolen data, identification of high-value victims, and optimization of ransom strategies at scale. This shift toward agentic AI allows attackers to automate reconnaissance, prioritize targets, and accelerate decision-making with minimal human intervention.

Security research indicates a sharp rise in attacks leveraging AI-enabled toolkits such as FraudGPT and WormGPT, which streamline phishing, malware development, and intrusion workflows. Ransomware ecosystems are also evolving in parallel, with operators such as FulcrumSec and INC Ransom adopting AI-driven data analysis techniques to sift through exfiltrated datasets, classify sensitive information, and maximize extortion pressure. This capability allows attackers to tailor ransom demands based on business impact, increasing both efficiency and success rates.

In 2026, threat actors have increasingly leveraged AI in their operations, enabling rapid analysis of massive volumes of stolen data, identification of high-value victims, and optimization of ransom strategies at scale.

AI-Enhanced Social Engineering

Social engineering campaigns increasingly feature highly interactive, AI-driven deception. As modern attack techniques increasingly involve multi-step, user-driven workflows that guide victims toward the manual execution of malicious actions, generative AI has amplified this threat by enabling hyper-personalized phishing campaigns. Threat actors can craft highly contextual and convincing messages using publicly available or previously compromised data, significantly increasing user engagement and initial access success rates. At the same time, AI-driven automation enables these campaigns to scale rapidly while maintaining personalization, making them harder to detect and defend against.

Additionally, a surge in impersonation attacks driven by deepfakes and voice cloning is reshaping identity-based threats. AI-generated audio, video, and text are increasingly being used to impersonate executives, IT staff, and trusted vendors with high levels of realism. These techniques are particularly effective in bypassing traditional verification controls, as they exploit human trust rather than technical vulnerabilities. As AI models continue to improve, similar campaigns are expected to become more frequent, scalable, and difficult to distinguish from legitimate communications.

A surge in impersonation attacks driven by deepfakes and voice cloning is reshaping identity-based threats.

Enterprise AI as an Emerging Attack Surface

As organizations rapidly adopt AI technologies, enterprise AI platforms are becoming attractive targets for threat actors. Prompt injection, in which attackers embed malicious instructions into inputs processed by AI systems, has emerged as a critical security risk. These hidden prompts can manipulate AI behavior, bypass safeguards, and trigger unintended actions, including data leakage or execution of unauthorized commands. Recent observations indicate a significant increase in prompt injection activity, with attackers experimenting to exploit AI-driven workflows. In some cases, hidden instructions embedded in emails or documents have successfully caused AI systems to exfiltrate sensitive data without user awareness. This ongoing evolution introduces an emerging attack vector in which AI systems act as intermediaries in data breaches.

Beyond data leakage, threat actors may attempt to poison AI-driven workflows by embedding persistent malicious logic within prompts. This can result in ongoing compromise, where AI systems unknowingly propagate malicious outputs or links over time. As a result, AI platforms are no longer just defensive tools but are increasingly becoming part of the enterprise attack surface.

AI platforms are no longer just defensive tools but are increasingly becoming part of the enterprise attack surface.

Security Implications and Strategic Considerations

The convergence of AI, social engineering, and ransomware operations represents a fundamental shift in the threat landscape. Adversaries are moving toward automated, intelligence-driven operations that combine data analysis, deception, and identity abuse to achieve scalable and efficient compromise.

To address these risks, organizations must treat AI inputs as untrusted, implement controls such as prompt validation and output monitoring, and ensure sensitive data is not directly exposed to AI systems without safeguards. In parallel, strengthening identity verification processes, enhancing user awareness against AI-driven deception, and monitoring for anomalous behavior across collaboration platforms remain critical. As AI continues to evolve, it will increasingly serve as both a force multiplier for threat actors and a new domain of risk, requiring organizations to rethink traditional security models and adopt proactive, intelligence-driven defense strategies.

Adversaries are moving toward automated, intelligence-driven operations that combine data analysis, deception, and identity abuse to achieve scalable and efficient compromise.

Outlook
for the remainder of 2026

So far in 2026, much of the threat landscape looks similar to the second half of 2025. Although its attack volume has decreased since its July 2025 peak, the Akira ransomware group will likely remain the most active threat in 2026, barring unforeseen law enforcement actions or an unexpected shuttering of the brand. Aside from Akira, established groups like Qilin and PLAY will likely continue to be responsible for the majority of incidents, though activity from any single group may fluctuate month to month.

Vulnerabilities in VPN and firewall products will likely remain a top initial access vector, and multiple threat groups continue to dedicate additional resources to bypassing security measures. Supply chain attacks are drawing increased attention from threat actors, particularly those focused on identity compromise involving API keys, tokens, and MFA bypass techniques rather than traditional credential theft. Threat groups also continue to see strong success with social engineering tactics, including the reemergence of Microsoft Teams–based campaigns first observed in early 2025, as well as device code phishing and ClickFix-style lures. Given the high likelihood that this trend will persist, user awareness and training are becoming increasingly critical, as security controls alone are often insufficient to fully mitigate these attack paths.

Lastly, while AI continues to play an expanding role in the cyber threat landscape, most current adversarial use of AI is generative rather than agentic, as threat actors are still experimenting with how AI fits into their operations rather than redesigning their operations around AI.

Akira will likely remain the most active threat in 2026, barring unforeseen law enforcement actions or an unexpected shuttering of the brand.

Vulnerabilities in VPN and firewall products will likely remain a top initial access vector.

Threat actors are still experimenting with how AI fits into their operations rather than redesigning their operations around AI.

As the threat landscape continues to evolve, Arete will continue to operationalize cyber intelligence to help organizations stay ahead of cyber threats and drive defensible decisions across response, remediation, recovery, and payment.

Appendix

Code Comparison, Lineage, and Evolution Of Conti, DragonForce, and Devman Ransomware

Following the 2022 leak of the Conti source code, multiple threat actors leveraged the exposed codebase to accelerate ransomware development. Rather than building frameworks from scratch, actors reused and modified existing implementations, leading to the emergence of several derivative families. Overall, the evolution of these ransomware families reflects a hybrid model of code reuse, builder-based development, and independent modification, rather than a strictly linear inheritance chain.

Although the DragonForce ransomware group initially used a variant of the leaked LockBit 3.0 builder in 2023, after launching their RaaS in 2024, strong alignment with Conti-derived implementations suggests that DragonForce likely transitioned to the leaked Conti-based codebase. Additionally, Devman, a former DragonForce affiliate, also used a modified version of DragonForce's ransomware, with the same similar overlaps to Conti's leaked source code.

High-Level Comparative Overview

FEATURE	CONTI	DRAGONFORCE	DEVMAN
Architecture	32-bit	32-bit	32/64-bit
Origin	Original	Early LockBit Black-linked, later Conti-based	DragonForce-based
Encryption Model	Hybrid	Hybrid	Hybrid
Symmetric Encryption	Salsa20 / ChaCha	Salsa20 / ChaCha	Salsa20 / ChaCha
Asymmetric Protection	RSA	RSA	RSA
Hashing	Murmur2	Murmur2	Murmur2
SMB Propagation	Yes	Yes	Yes
Tor Infrastructure	Yes	Yes	Yes
Mutex	hsfjuukjzloqu28oajh727190	hsfjuukjzloqu28oajh727190	hsfjuukjzloqu28oajh727190 (32-bit)
Ransom Note Name	readme.txt	readme.txt	readme.txt / RESTORE_FILES.txt

Shared Cryptographic Implementation (ChaCha Core)

Conti, DragonForce, and Devman all use an identical ChaCha-like ARX encryption routine, indicating shared or reused cryptographic code.

```
.text:00405DAD                  add     edx, ebx
.text:00405DAF                  mov     [ebp+var_84], ebx
.text:00405DB5                  mov     [ebp+var_50], edx
.text:00405DB8                  xor     edx, eax
.text:00405DBA                  mov     eax, [ebp+var_64]
.text:00405DBD                  add     eax, esi
.text:00405DBF                  rol     edx, 7
.text:00405DC2                  mov     [ebp+var_64], eax
.text:00405DC5                  mov     [ebp+var_7C], edx
.text:00405DC8                  mov     edx, [ebp+var_60]
.text:00405DCB                  xor     edx, eax
.text:00405DCD                  mov     eax, [ebp+var_58]
.text:00405DD0                  rol     edx, 10h
.text:00405DD3                  add     eax, edx
.text:00405DD5                  mov     [ebp+var_58], eax
.text:00405DD8                  xor     eax, esi
.text:00405DDA                  mov     esi, [ebp+var_64]
.text:00405DDD                  rol     eax, 0Ch
.text:00405DE0                  add     esi, eax
.text:00405DE2                  xor     edx, esi
.text:00405DE4                  mov     [ebp+var_64], esi
.text:00405DE7                  mov     esi, [ebp+var_58]
.text:00405DEA                  rol     edx, 8
.text:00405DED                  add     esi, edx
.text:00405DEF                  mov     ebx, esi
.text:00405DF1                  mov     [ebp+var_58], esi
.text:00405DF4                  mov     esi, [ebp+var_90]
.text:00405DFA                  xor     ebx, eax
.text:00405DFC                  mov     eax, [ebp+var_54]
.text:00405DFF                  add     eax, edi
.text:00405E01                  rol     ebx, 7
.text:00405E04                  xor     esi, eax
```

Figure 3 | Shared Cryptographic Implementation (Source: Arete)

Command Line Arguments

Conti, DragonForce, and early Devman (32-bit) share identical command-line arguments:

COMMAND LINE PARAMETERS	DESCRIPTION
-p	Path to encrypt file
-m	EncryptMode
-log	Create log file
-size	File encryption percentage
-nomutex	Skip mutex check

Devman Divergence

However, newer Devman variants (notably 64-bit / Rust-based samples) introduce a completely different command-line interface:

COMMAND LINE PARAMETERS	DESCRIPTION
-h, --help	Show this help
-v, --verbose	Show every file
-p, --path <dir>	Encrypt specific path only
--stealth	Self-delete after
--debug	Debug mode (no evasion)

Conti vs. DragonForce

DragonForce demonstrates overlapping characteristics with Conti-derived ransomware families. However, these similarities are not consistent across all samples, indicating a multi-stage evolution rather than direct inheritance from a single codebase.

Key Shared Traits:

COMPONENT	Conti	DRAGONFORCE
Binary architecture	32-bit	32-bit
Encryption	Salsa20 / ChaCha	Salsa20 / ChaCha
Hashing	Murmur2	Murmur2

Other key capabilities found in DragonForce and Conti:

- Obfuscated stack strings
- XOR data encoding
- PEB access
- API resolution via PE exports
- PE header & sections parsing
- SMB lateral movement
- Tor-based DLS

Shared Mutex Artifact:

The same ***"hsfjuukjzloqu28oajh727190"*** mutex is used across Conti, DragonForce, and Devman (32-bit) samples.

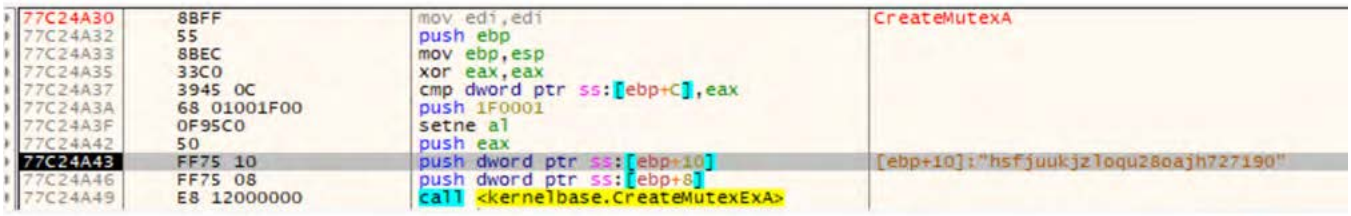

Figure 4 | Mutex Value of Conti (Source: Arete)

Figure 5 | Mutex Value of DragonForce (Source: Arete)

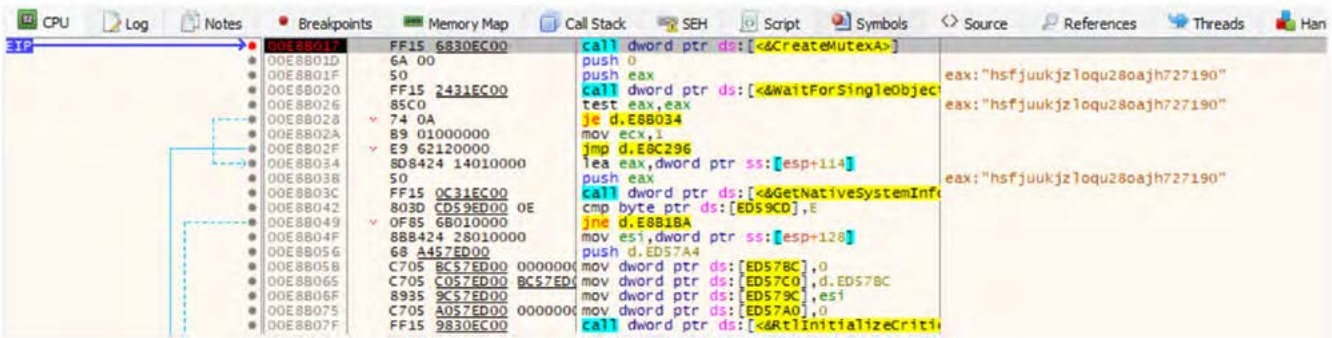

Figure 6 | Mutex Value of Devman [32-bit] (Source: Arete)

Older DragonForce samples, however, uses unique, GUID-like mutex value, such as: ***Global\\5fc8a5eb49ff2e025305c42026781307***

Figure 7 | Mutex Value of Older DragonForce (Source: Arete)

DragonForce vs. Devman

Older 32-bit Devman ransomware samples (VC++8 compiled) are functionally and structurally identical to DragonForce.

Shared Components are:

- Command line parameters
- Mutex value
- Process and service kill lists
- File and folder exclusions
- Encryption workflow
- Decryption logic

Devman 32-bit vs. DragonForce 32-bit Ransomware

Devman sample hash (32-bit)	DragonForce sample hash (32-bit)	Bindiff similarity score
df5ab9015833023a03f92a797e20196672c1d6525501a9f9a94a45b0904c7403	b10129c175c007148dd4f5aff4d7fb61eb3e4b0ed4897fea6b33e90555f2b845	99%

Figure 8 | BinDiff results of DragonForce and Devman (Source: Arete)

Footnote

Data Collection and Analysis Methodology

Arete provides comprehensive incident response services, and the insights shared in this report are derived from incident response, digital forensics, restoration, managed detection and response, endpoint protection, threat hunting, threat intelligence, threat actor communications, dark web monitoring, and advisory and consulting services. While not every client opts to use all the cyber solutions Arete offers, Arete gathers data points from thousands of unique ransomware engagements going back to 2018. By collecting and validating data from diverse sources, Arete builds a comprehensive threat intelligence repository, analyzes raw data, identifies patterns, and provides context to enable informed decision-making.

All data pertaining to threat actors is collected and analyzed to ensure victims are anonymized and there is no chance of threat actors or readers identifying any victim. Only data from incidents where victims were extorted by the threat actor, with or without encryption, are included in this report. While we share some insights from pre-ransomware attacks in which threat actors were disrupted prior to encrypting and/or stealing data, those incidents are not included in any statistics. Finally, any information that Arete assesses could be used by threat actors to improve their operations (e.g., negotiated discounts per threat actor) is excluded from public reports but available to trusted partners upon request.

Bias Acknowledgement

There are thousands of ransomware attacks claimed by threat actors worldwide each year, while many more likely go unreported or remain unknown to the victims. Arete conducts analysis based on the data collected during our incident response engagements. These incident response engagements primarily represent organizations who have cyber insurance. As our data represents just a sample of the overall number of global ransomware attacks, it creates a sampling bias. The analysis contained in this report reflects the trends Arete observes first-hand during our engagements with cybercriminals and may differ from trends observed by the greater cyber community.

Sources

Arete Internal Data

Ransomware.live. (n.d.). https://www.ransomware.live

RansomLook — **Open ransomware intelligence.** (n.d.). https://www.ransomlook.io/

Threat Actor Profile: DragonForce Ransomware Group. (2025, February 20). Cyble. https://cyble.com/threat-actor-profiles/dragonforce-ransomware-group/

DragonForce Ransomware. (2026). **In Black Point Cyber**. https://blackpointcyber.com/wp-content/uploads/2026/02/DragonForce-1.pdf

Kichatov, N., Low, S., Kashtanov., (2025, December 29). **Inside the Dragon: DragonForce Ransomware Group.** Group-IB. https://www.group-ib.com/blog/dragonforce-ransomware/

CyberMasterV. (2021b, July 6). **Dissecting the last version of Conti Ransomware using a step-by-step approach.** CYBER GEEKS. https://cybergeeks.tech/dissecting-the-last-version-of-conti-ransomware-using-a-step-by-step-approach/

Virtusio, D., Catalan Alegre, D., Kimhy, E., & Pontiroli, S. (2025, November 4). **The DragonForce Cartel: Scattered Spider at the gate.** Acronis. https://www.acronis.com/en/tru/posts/the-dragonforce-cartel-scattered-spider-at-the-gate/

TrendAI™. (2025, October 29). **Ransomware Spotlight Dragonforce.** TrendAI™. https://www.trendaisecurity.com/en-us/resources-insights/research/ransomware-spotlight-dragonforce

DiMaggio, J. (2025, October 15). **Devman's RaaS Launch: The Affiliate Who Aims to Become the Boss.** Analyst1. https://analyst1.com/devmans-raas-launch-the-affiliate-who-aims-to-become-the-boss/

ANY.RUN. (2025, December 19). **DEVMAN: Ransomware Overview.** Medium. https://medium.com/@anyrun/devman-ransomware-overview-32600d4da684

(2026, January 26). **Dark Web Profile: BravoX.** SOCRadar®. https://socradar.io/blog/dark-web-profile-bravox-ransomware/

Carvey, H., & O'Donnell-Welch, L. (2026, April 7). **Leveling Up with NightSpire Ransomware.** Huntress. https://www.huntress.com/blog/nightspire-ransomware

Pompeu, L. (2026, February 25). **NightSpire Ransomware: How it works and how to defend against it.** https://www.provendata.com/blog/nightspire-ransomware

Verma, R. (2025, September 4). **NightSpire Ransomware Expands Reach with Aggressive Extortion Deadlines.** Hive Pro. https://hivepro.com/threat-advisory/nightspire-ransomware-expands-reach-with-aggressive-extortion-deadlines/

Ilascu, I. (2026, May 11). **Instructure confirms hackers used Canvas flaw to deface portals.** BleepingComputer. https://www.bleepingcomputer.com/news/security/instructure-confirms-hackers-used-canvas-flaw-to-deface-portals/

(2026, March 10). FortiGate Edge intrusions | **Stolen service accounts lead to rogue workstations and deep AD compromise.** SentinelOne. https://www.sentinelone.com/blog/fortigate-edge-intrusions/

Arghire, I. (2025, December 22). **WatchGuard patches Firebox Zero-Day exploited in the Wild.** SecurityWeek. https://www.securityweek.com/watchguard-patches-firebox-zero-day-exploited-in-the-wild/

(2026, March 31). **WatchGuard Fireware OS VPN Vulnerability | Active Exploitation.** Horizon3.ai. https://horizon3.ai/attack-research/vulnerabilities/cve-2025-14733/

(2026, February 2). **Advisory ID: BT26-02.** BeyondTrust. https://www.beyondtrust.com/trust-center/security-advisories/bt26-02

Moore, J. (2026, February 19). **VSHELL and SparkRAT observed in exploitation of BeyondTrust critical vulnerability (CVE-2026-1731).** Unit 42. https://unit42.paloaltonetworks.com/beyondtrust-cve-2026-1731/

Security Arsenal Team. (2026, March 16). **Defending against Microsoft Teams Social Engineering: Blocking IT impersonation attacks.** Security Arsenal. https://securityarsenal.com/blog/defending-against-microsoft-teams-social-engineering-blocking-it-impersonation-attacks

Microsoft Defender Security Research Team & Microsoft Defender Experts (2026, April 18). **Cross-tenant helpdesk impersonation to data exfiltration: A human-operated intrusion playbook.** Microsoft. https://www.microsoft.com/en-us/security/blog/2026/04/18/crosstenant-helpdesk-impersonation-data-exfiltration-human-operated-intrusion-playbook/

Stone-Gross, B. (2026, April 16). **Payouts King takes aim at the ransomware throne.** ZScaler™. https://www.zscaler.com/blogs/security-research/payouts-king-takes-aim-ransomware-throne

Louw, J. (2026, March 6). **InstallFix: How attackers are weaponizing malvertised install guides.** Push Security. https://pushsecurity.com/blog/installfix

Bitam, S. (2026, February 19). **MIMICRAT: ClickFix Campaign Delivers Custom RAT via Compromised Legitimate Websites.** Elastic Security Labs. https://www.elastic.co/security-labs/mimicrat-custom-rat-mimics-c2-frameworks

Insikt Group®. (2026, March 25). **ClickFix campaigns targeting Windows and MacOS.** Recorded Future. https://www.recordedfuture.com/research/clickfix-campaigns-targeting-windows-and-macos

Arntz, P. (2026, March 30). **New macOS security feature will alert users about possible ClickFix attacks.** Malwarebytes™. https://www.malwarebytes.com/blog/news/2026/03/new-macos-security-feature-will-alert-users-about-possible-clickfix-attacks

Microsoft Defender Security Research Team. (2026, February 5). **New Clickfix variant 'CrashFix' deploying Python Remote Access Trojan.** Microsoft. https://www.microsoft.com/en-us/security/blog/2026/02/05/clickfix-variant-crashfix-deploying-python-rat-trojan/

Dutta, T. S. (2025, November 19). **Destructive Akira Ransomware Attack with a Single Click on CAPTCHA in Malicious Website.** Cyber Security News. https://cybersecuritynews.com/destructive-akira-ransomware-attack/

The Fortinet 2026 Global Threat Landscape report reveals a surge in AI-Enabled cybercrime, contributing to a 389% increase in ransomware victims Year-over-Year. (2026, April 30). Fortinet. https://investor.fortinet.com/news-releases/news-release-details/fortinet-2026-global-threat-landscape-report-reveals-surge-ai

Shapland, R. (2026, January 5). **Real-world AI voice cloning attack: A red teaming case study.** TechTarget | Search Security. https://www.techtarget.com/searchsecurity/tip/Real-world-AI-voice-cloning-attack-A-red-teaming-case-study

Kobie, N. (2026, May 5). **Five Eyes agencies sound alarm over risky agentic AI deployments.** ITPro. https://www.itpro.com/security/five-eyes-agencies-sound-alarm-over-risky-agentic-ai-deployments

ByteVanguard. **Prompt Injection: The New Phishing Layer in Enterprise AI.** (2026, March 10). ByteVanguard. https://bytevanguard.com/2026/03/10/prompt-injection-enterprise-ai/

General Inquiries

info@areteir.com
646-907-9767

For Immediate Assistance with a Cyber Incident

USA & Canada

arete911@areteir.com
866-210-0955

India

arete112@areteir.com
1800-890-7383

UK

arete999@areteir.com
+44-800-208-8084