Q2 2026

ThreatCluster

Threat Intelligence Report

Published: July 2026
Classification: TLP:CLEAR

Executive summary

ThreatCluster processed 85,433 articles from 10,785 sources in Q2 2026, up 56% and 57% on Q1, and grouped them into 6,714 clusters. Since the platform's first full quarter of operation in late 2025, article intake has nearly tripled while the cluster count rose 39%, so articles per cluster climbed from 6.0 to 12.7. The collection is widening faster than the set of distinct stories, so mention counts increasingly measure how much attention each story drew.

16,414 sources
170,286 articles
17,668 clusters
Total platform numbers

89.6%
noise reduction

Across three quarters of the platform's own data, the defining pattern is fragmentation. Coverage of the top five ransomware brands fell from 37% of all ransomware mentions at inception to 21% this quarter, while the number of distinct operations drawing coverage rose to 183.

Distinct malware families more than doubled since inception, from 507 to 1,174. Distinct nation-state groups rose from 246 to 583. Attention is spreading across more actors, not concentrating behind fewer, and no single ransomware brand has replaced the operations dismantled over the past two years. Qilin stayed top but flat; the climbers were newer or rebranded operations, chiefly World Leaks and The Gentlemen.

Three storylines sat above the collection growth. Supply-chain compromise rose 85% as an attack type and 194% as a technique, driven by the Mini Shai-Hulud npm worm run by TeamPCP.

ThreatCluster
One Feed. Every Threat.
threatcluster.io
hello@threatcluster.io
71-75 Shelton Street, Covent Garden, London, WC2H 9JQ

Coverage of the financial (up 130%), energy (up 115%) and defence (up 185%) sectors rose well ahead of the 56% growth in intake. And the quarter's geography tracked live conflict: Russia, Ukraine, the United States, Iran and China led, with a Middle East cluster around the Iran and Israel military events and an India and Pakistan cluster around Operation Sindoor.

AI ran through the quarter on three tracks.

- As a stolen asset, Anthropic accused operators tied to Alibaba's Qwen lab of the largest model-distillation campaign it has disclosed.
- As a weapon, an open-weight model matched a frontier system on an offensive-security benchmark and stolen agent installs were used to breach companies in the wild.
- As a target, prompt-injection coverage rose 367% and an AI-targeting malware category took shape. The counts throughout measure reporting attention rather than incidence; everything below is read on that basis.

The quarter in numbers

ThreatCluster aggregates security reporting from thousands of open sources, groups it into clusters using density-based semantic clustering, then tags and counts the entities in each story. A mention is the number of distinct articles in which an entity was tagged, merged to canonical names across known aliases. Where a hard figure appears, a victim total, a CVSS score, a patch date, it comes from a named primary source and is cited inline.

Comparisons run across three periods: a baseline of Q4 2025 (November to December, the first full quarter of platform operation, with a nine-day October stub excluded), Q1 2026, and Q2 2026.

Intake rose 56% between Q1 and Q2, so a quarter-on-quarter rise above roughly that line reflects real movement in attention while a smaller one mostly reflects wider collection. Baseline multiples for MITRE techniques and the largest country tallies over-read, because canonical tagging was still rolling out in late 2025; for those, the quarter-on-quarter figure is the reliable trend. Percentages off very low counts are flagged as small-base.

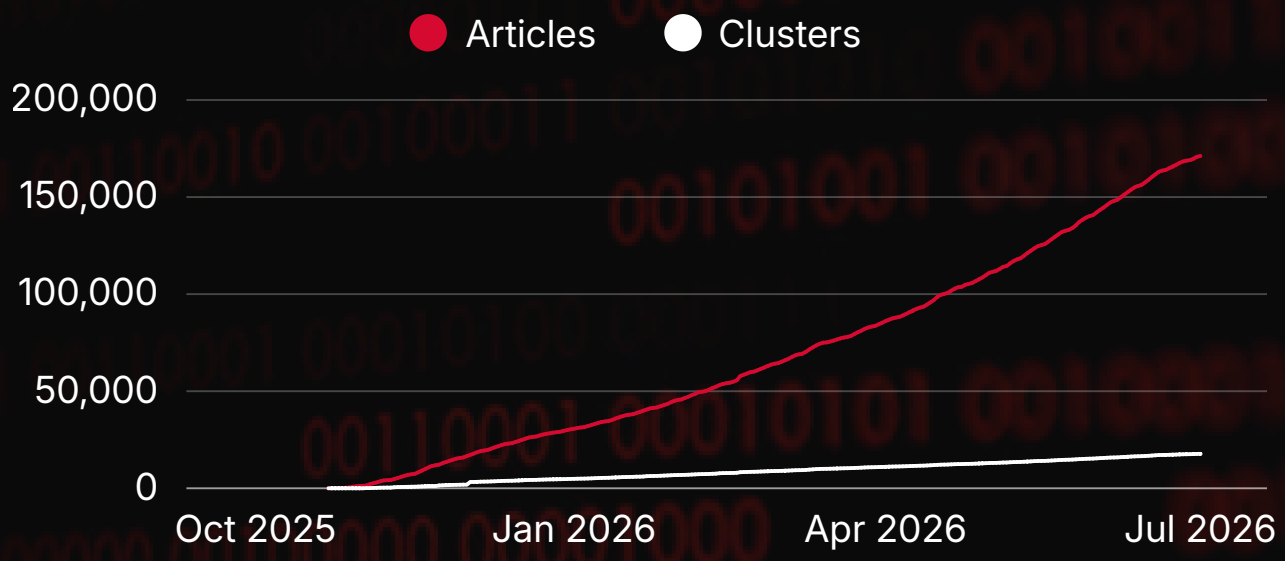

The source base has grown faster than the story count in every period. That is the mechanism behind most of the raw rises in this report, and it is also the reason the concentration figures below matter more than any single count: they measure how attention is distributed, independent of how much was collected.

Metric	Baseline (Q4 2025)	Q1 2026	Q2 2026	Q2 vs Q1
Articles	29,196	54,861	85,433	+56%
Sources	4,173	6,878	10,785	+57%
Clusters	4,841	6,151	6,714	+9%
Articles per cluster	6	8.9	12.7	+43%

The clearest pattern is fragmentation. Across ransomware, malware and nation-state groups, the number of distinct names drawing coverage keeps climbing while the five most-covered names hold a shrinking or flat share, so attention is spreading across more players rather than concentrating behind a few. Comparing the platform's first full quarter (Q4 2025) with Q2 2026:

Category	Coverage share held by the top 5	Distinct names covered	Category	Coverage share held by the top 5
Ransomware groups	37% to 21%	165 to 183	Ransomware groups	37% to 21%
Malware families	10% to 7%	507 to 1,174	Malware families	10% to 7%
Nation-state groups	14% to 18%	246 to 583	Nation-state groups	14% to 18%

Ransomware and malware are fragmenting: the leading names hold a smaller share as the field grows, with distinct malware families more than doubling. Nation-state tracking is broadening: distinct groups more than doubled while the top five held their share. Industries, countries and attack types stayed concentrated and broadly flat over the same period.

How the quarter unfolded

April

The quarter opened on the US and Iran military events, with Operation Epic Fury and Operation Midnight Hammer coverage peaking and APT28 hitting its quarterly high of 95 mentions in a single month.

North Korean operators set the tone in crypto: the $285 million Drift exchange hack on 1 April was followed on 18 April by the Kelp DAO bridge exploit, at $293 million the largest DeFi theft of the year, attributed by security researchers to Lazarus Group.

BlackCat coverage spiked around US affiliate and insider prosecutions, and two operations that would run through the quarter first appeared on 7 April: the Silent Ransom Group and Anthropic's Project Glasswing.

May

May was the exploitation month. Mini Shai-Hulud hit TanStack's build pipeline on 11 May, its operators published the source code on 12 May, and the AntV wave followed on 19 May; supply-chain technique coverage peaked at 201 mentions.

Cisco disclosed the SD-WAN zero-day CVE-2026-20182 on 14 May under an emergency patching directive.

The Canvas breach broke at the start of the month and ended with Instructure paying ShinyHunters' ransom on 11 May, and the same group began exploiting the Oracle PeopleSoft zero-day on 27 May.

Operation Sindoor peaked at 30 mentions as the India and Pakistan coverage crested, and Check Point published its analysis of The Gentlemen's leaked backend.

June

June belonged to extortion and enforcement.

World Leaks took 28 of its 29 quarterly mentions in a single month, Operation Endgame surged to 28 as the takedown campaign returned, and FortiBleed emerged mid-month with a CISA alert on 18 June.

The AI story compressed into one week: Anthropic's Senate letter on the Qwen distillation campaign on 10 June, US Commerce restrictions on its own models two days later, and the GLM-5.2 open-weight release the day after that.

StealC spiked to 32 mentions, data-breach coverage hit its monthly high of 1,048 articles, and the China-nexus cluster UNC6508 appeared from nothing with 13 mentions.

The quarter's biggest stories

Clusters are the platform's unit of story, so the largest clusters are the quarter's biggest stories by volume of coverage.

Anthropic's Claude Mythos Preview Sparks Cybersecurity Revolution	**1,126** Articles	**69.9** Threat Score
Mini Shai-Hulud Supply Chain Attack Targets SAP npm Packages	**601** Articles	**70.5** Threat Score

threatcluster.io
hello@threatcluster.io
71-75 Shelton Street, Covent Garden, London, WC2H 9JQ

Major Data Breach at Canvas Affects 275 Million Users	**590** Articles	**71.0** Threat Score
Moscow Issues Threats to Baltic States Over Ukrainian Drone Strikes	**265** Articles	**59.9** Threat Score
Rise of AI-Powered Cyber Attacks: New Threat Landscape Emerges	**195** Articles	**69.5** Threat Score

Volume is one cut; severity is the other, and they barely overlap.

Critical Joomla JCE Vulnerability Under Active Exploitation	**11** Articles	**87.2** Threat Score
Critical Zero-Day Vulnerability CVE-2026-20182 Exploited in Cisco SD-WAN Systems	**131** Articles	**87.2** Threat Score
Critical WebLogic RCE Vulnerability Exploited in the Wild	**2** Articles	**86.2** Threat Score

Critical Unauthenticated RCE Vulnerability in LiteLLM Exploited in the Wild	**17** Articles	**86.0** Threat Score
Critical Zero-Day Vulnerability in LiteSpeed cPanel Plugin Actively Exploited	**17** Articles	**86.0** Threat Score

The two lists share exactly one entry. The Cisco SD-WAN cluster was the only story of the quarter that was both heavily covered (131 articles) and top-of-scale severe (87.2); everything else splits cleanly.

The volume list is AI, mega-breaches and geopolitics; the severity list is active exploitation, and most of its clusters drew fewer than 25 articles, with a WebLogic RCE exploited in the wild covered by just 2.

That gap is the argument for severity scoring: the most dangerous stories of the quarter were mostly not the most covered ones.

Two of the severity clusters also connect to threads elsewhere in this report: the LiteLLM RCE is part of TeamPCP's AI supply-chain activity, and the GRU home-router cluster sits behind the group's 1,200% rise in the APT table.

Ransomware Activity

Qilin held the top spot across all three periods without growing, and the top-five share of ransomware coverage has fallen every quarter since inception. The clearest signal is the rise of newer and rebranded operations alongside a set of long-defunct brands resurfacing for reasons that have nothing to do with live operations.

Group	Baseline	Q1	Q2	Q2 vs Q1
Qilin	53	50	52	+4%
BlackCat/ALPHV	50	8	37	+362%
LockBit	37	32	35	+9%
Akira	69	20	30	+50%
World Leaks	0	2	29	+1350%
Conti	12	15	26	+73%
The Gentlemen	8	1	24	+2300%
Nitrogen	0	7	23	+229%
Clop	106	23	22	-4%
Black Basta	3	22	20	-9%
Rhysida	10	7	20	+186%
REvil	3	5	19	+280%

Behind the coverage counts, ThreatCluster's dark web monitoring gives a second view of the same groups. As of this report it tracks 120 ransomware and extortion operations across 475 underground sites, with 8,056 victims posted to leak sites, and 1,014 known breaches recorded to date.

120
Ransomware Groups

8,056
Total Victims

475
Underground Sites

1,014
Recorded Breaches

World Leaks rose from 2 clustered articles in Q1 to 29, with 28 of those landing in June, the sharpest single-month surge in the table.

It is the extortion-only successor to Hunters International, itself a likely Hive descendant, and it abandoned encryption for pure data theft and publication, marketed to affiliates as an Extortion-as-a-Service platform with a custom exfiltration tool and an "Insider" portal that gives selected media 24-hour early access to stolen data (Group-IB; Halcyon).

ThreatCluster's monitoring of the group's leak site shows postings concentrated in healthcare, manufacturing and business services, with US organisations the most numerous. Epsilon Research investigated one of those victims, Tata Electronics, from the group's leak data this year. Darktrace noted that despite the encryption-free billing, at least one 2026 incident it handled involved both exfiltration and an encryption payload, so the model is not as clean in practice as the marketing.

The Gentlemen rose from a single clustered article in Q1 to 24, accelerating through the quarter (4 in April, 5 in May, 15 in June), and the underlying activity is larger than the coverage suggests.

Check Point counted more than 320 public victims in 2026, The Hacker News reported the group's own claim of 478, and NTT ranked it the second most active operator after Qilin, at around 10% of observed attacks.

Microsoft tracks the operation as Storm-2697, a self-propagating, worm-like ransomware run by a former Qilin affiliate on a 90/10 split, with access drawn mainly from stolen credentials and exposed edge devices.

The most-covered episode was Check Point's May analysis of the group's own leaked backend, which showed it had breached a UK software consultancy in April and reused that access to hit one of the firm's clients in Turkey, naming the consultancy on its leak site as the "access broker". Check Point also reported the group uses Chinese AI models, including DeepSeek and Qwen, to build its ransomware, which links this story to the open-weight-model coverage later in the report.

BlackCat/ALPHV's 362% quarter-on-quarter rise is a court story, not a return to operations, and coverage still sits 26% below its Q4 2025 baseline.

Its Q2 mentions were front-loaded in April (25 of 37) around a run of US affiliate and insider prosecutions, including a former DigitalMint ransomware negotiator who pleaded guilty to passing victim information to the group (Help Net Security).

The same reading applies to Conti, REvil, WannaCry and NotPetya, long-defunct names whose presence in coverage reflects leaked-source reuse, alumni, decryptor releases and analogy rather than live operations.

Silent Ransom Group is the cleanest genuinely new entrant, first seen on 7 April and reaching 12 mentions on a rising monthly path. Also tracked as Luna Moth, Chatty Spider and UNC3753, it is an extortion-without-encryption crew that the FBI has flagged for callback-phishing campaigns against US law firms, and which Resecurity tied to DNS fast-flux infrastructure.

Ransomware as an attack type fell 14% quarter on quarter even as the number of distinct brands grew.

Coverage is shifting from encryption ransomware toward data-theft extortion, which surfaces under data breach and supply-chain labels rather than under the ransomware tag. That makes backup-centred defence necessary but insufficient, and puts the weight on exfiltration detection and credential hygiene.

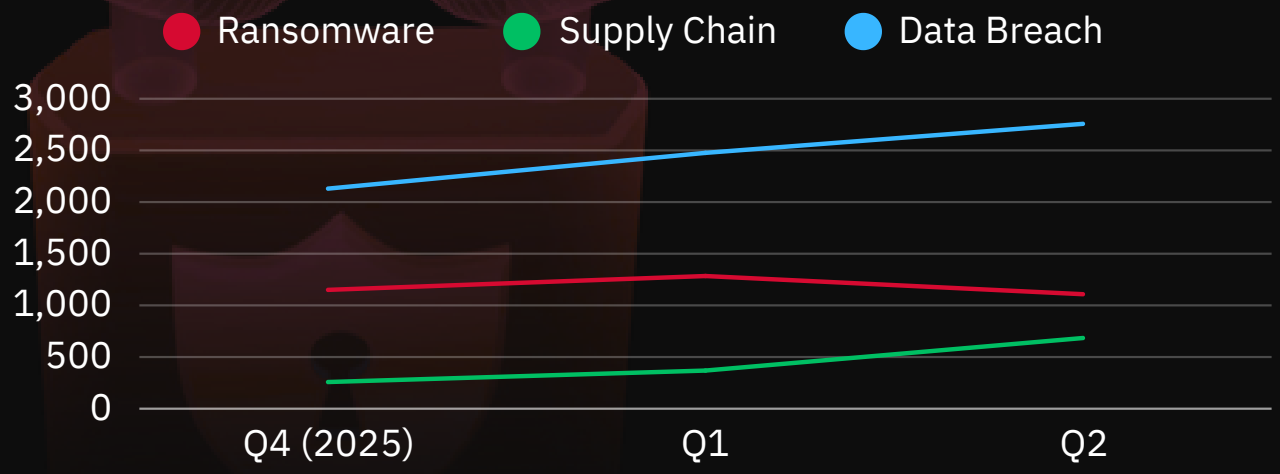

Coverage versus leak-site activity

Because the platform holds both sides, media mentions and leak-site postings can be set against each other for the same groups. The ratio of posted victims to articles shows which operations run below the media radar.

Group	Q2 articles	Q2 leak-site victims	Victims per article
Qilin	47	213	5
The Gentlemen	18	193	11
DragonForce	16	132	8

Group	Q2 articles	Q2 leak-site victims	Victims per article
Akira	30	102	3
INC Ransom	1	73	73
NightSpire	3	64	21.3
Nova	2	59	29.5
Krybit	2	50	25
SafePay	3	43	14.3
Play	5	34	6.8
Stormous	0	26	no coverage
World Leaks	30	21	0.7

The mismatch runs in both directions. Qilin posted the most victims and drew the most coverage, yet still posted 4.5 victims for every article written about it.

The Gentlemen's 193 postings against 18 articles corroborate the external counts cited earlier: the second most active operation of the year was running an order of magnitude ahead of its press.

The middle of the table is the quiet zone: INC Ransom posted 73 victims and drew one article, and NightSpire, Nova, Krybit and SafePay each posted between 43 and 64 victims on two or three articles apiece.

Below them, another four brands with between 22 and 54 posted victims drew no coverage at all in the quarter. The single inversion is World Leaks, the only group with more articles than victims, at 0.7: coverage follows a novel extortion model, not victim volume. Media attention is a brand-recognition market, and the leak-site telemetry is the corrective.

Nation-state and APT activity

Distinct nation-state groups drawing coverage rose from 246 at inception to 583 this quarter, with roughly 377 genuinely new after de-duplication, mostly UNC-numbered clusters and vendor cryptonyms. The top five held a flat share, so this is breadth rather than a shift in who dominates.

Lazarus Group led at 167 mentions, up 241% quarter on quarter, consistent with continued North Korea-aligned activity against developers and cryptocurrency targets.

The quarter's largest single theft anchored that coverage: the 18 April Kelp DAO bridge exploit drained $293 million, the biggest DeFi theft of the year, attributed by security researchers to Lazarus and its TraderTraitor subunit, and it came just weeks after the $285 million Drift exchange hack on 1 April was also linked to North Korean actors.

TraderTraitor itself entered the table new at 18 mentions. North Korea featured in 706 articles overall, up 228%, and the quarter's reporting included the macOS.Gaslight implant analysed by SentinelLABS, which embeds 38 fabricated system messages to derail AI-assisted malware triage, and the Proofpoint UNK_DeadDrop campaign luring developers with fake recruiter personas and malicious npm and PyPI packages. North Korean breadth is also the clearest illustration of the naming trend behind the 583 distinct groups: coverage that would once have been filed under Lazarus alone now splits across at least seven DPRK-linked names in the data.

Kimsuky held steady at 26, though its coverage collapsed to a single mention in June; APT37 rose 86% to 13; APT45 took all 15 of its mentions in May; and the crypto-theft reporting spread across TraderTraitor (18), Sapphire Sleet (14, back from two prior mentions) and UNC1069 (19, up 90%).

Lazarus coverage itself crested at 88 mentions in May as the Kelp DAO laundering and asset-freeze story ran on, then fell to 31 in June. Contagious Interview, the recruitment-lure operation behind much of the developer targeting, held around its baseline at 10 mentions.

Group	Baseline	Q1	Q2	Q2 vs Q1
Lazarus Group	40	49	167	2.41
APT28	22	91	139	0.53
ShinyHunters	14	31	43	0.39
Volt Typhoon	16	21	36	0.71
Sandworm	15	24	35	0.46
APT29	2	6	32	4.33
TeamPCP	0	7	27	2.86
GRU	7	2	26	12
Kimsuky	24	26	26	0%
Salt Typhoon	16	46	25	-46%
Scattered Spider	24	9	24	1.67
APT41	5	13	21	0.62

The Russian cluster is the largest geopolitical signal. APT28 rose 53% to 139 mentions but was front-loaded in April (95 of its Q2 mentions, falling to 9 by June), Sandworm rose 46% to 35 around wiper activity against Ukrainian organisations, and APT29 re-emerged from 6 mentions to 32, a 433% rise on a steady monthly climb. ThreatCluster's country counts tie this to the war:

Russia appeared in 2,260 articles (up 245%) and Ukraine in 1,766 (up 328%), the two largest country movements in the dataset.

Among the Chinese clusters the two Typhoons diverged: Volt Typhoon coverage rose 71% to 36 while Salt Typhoon's fell 46% to 25, and UNC3886, prominent in Q1 at 48, dropped to 14. The quarter also brought the Cisco SD-WAN exploitation covered below, which Cisco tied to an actor it calls UAT-8616 whose infrastructure overlaps China-nexus operational relay networks.

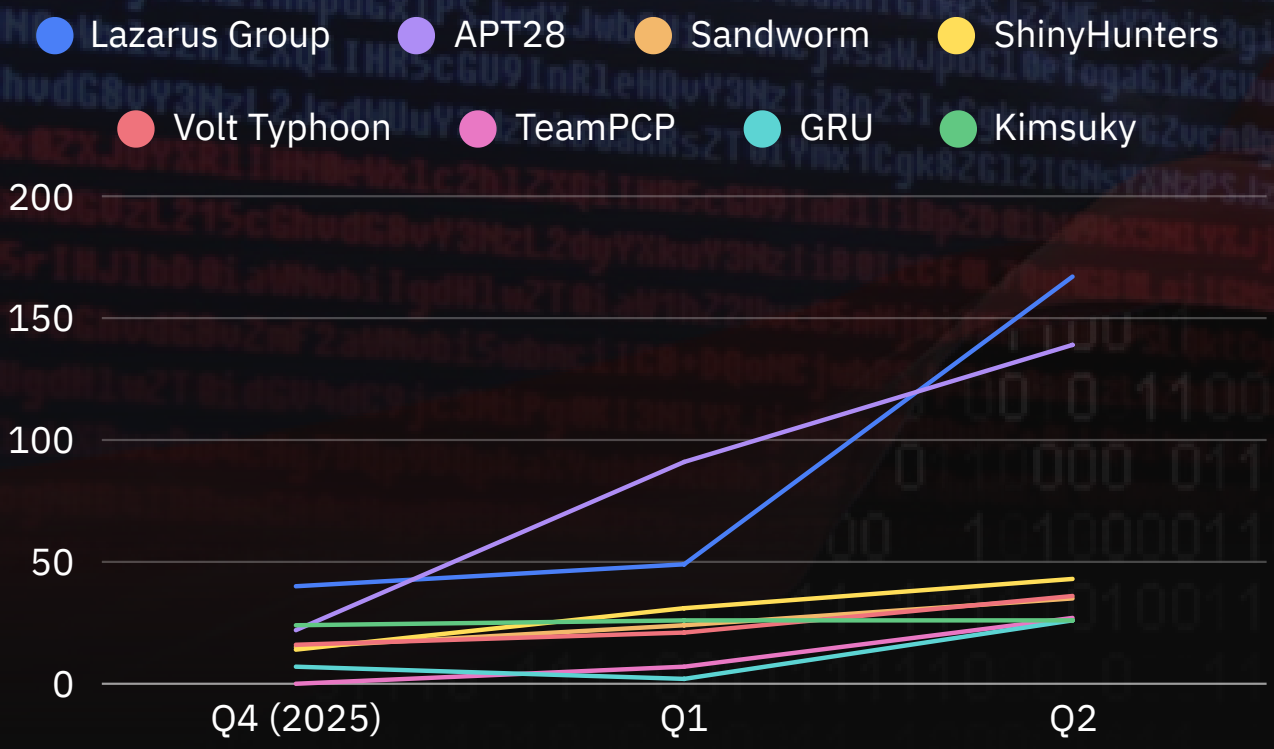

TeamPCP (also tracked as UNC6780) is the entry to watch, up from zero at baseline to 27 mentions; its real footprint is in the supply chain, where it ran the Mini Shai-Hulud npm worm and the LiteLLM and Trivy compromises.

ShinyHunters rose 39% to 43 and was tied by reporting to exploitation of the Oracle PeopleSoft zero-day CVE-2026-35273; it also appears in the ransomware data under its BreachForums extortion persona.

Exploitation and zero-days

Zero-day exploitation held near the top of the attack-type table, rising 11% quarter on quarter to 1,027 mentions and 66% above the Q4 2025 baseline. The reliable signal is which exploitation stories dominated, and the quarter belonged to edge and management-plane appliances.

Three events carried most of it.

Scattered Spider's 167% rebound fits its pattern of social-engineering-led intrusions against retail and insurance. The clearest new arrivals were UNC6508, a China-nexus cluster that appeared only in June, and the Nimbus Manticore, VerdantBamboo and UNC1151 clusters.

The largest was the Cisco Catalyst SD-WAN chain: a CVSS 10 authentication-bypass zero-day, CVE-2026-20182, disclosed on 14 May, which Cisco attributed to the China-nexus actor UAT-8616 and which drew a CISA emergency directive to patch within days; it was one of 7 Cisco SD-WAN zero-days reported across the year.

In late May and June, ShinyHunters exploited an Oracle PeopleSoft zero-day (CVE-2026-35273) that Mandiant tied to breaches at more than 100 organisations, most of them universities, which drove much of the education-sector rise.

In mid-June, FortiBleed exposed cracked administrator credentials for tens of thousands of internet-facing Fortinet firewalls across 194 countries, prompting a CISA alert.

The pattern held across the quarter: exposed management interfaces and widely deployed appliances drew the heaviest exploitation coverage, and public proof-of-concept code reliably preceded a wave of opportunistic exploitation. In the technique data, exploitation of public-facing applications (T1190) accelerated through the quarter (77 mentions in April, 90 in May, 142 in June) and web-shell activity (T1505.003) rose 348% quarter on quarter.

Most-covered CVEs

Two things stand out. First, the list skews to Linux kernel and distribution advisories, which is a property of the source mix as much as the threat picture: the platform's largest single source is a Linux security advisory outlet, and advisory coverage cites identifiers densely.

Second, the counts are conservative by design. A CVE is counted only when the identifier itself is tagged in an article, so flaws that dominated coverage under a product name rather than a number sit below their real footprint; the Cisco SD-WAN zero-day's cluster carried 131 articles while the CVE identifier drew fewer mentions than anything in this table.

The active-exploitation entries are the ones to hold: the PeopleSoft zero-day, the FortiGate exploitation and the F5 NGINX pair all carried the quarter's highest severity scores in this set.

CVE	Mentions	Associated story
CVE-2026-31431	61	Moldova healthcare breach
CVE-2026-43284	50	90-day disclosure policy debate
CVE-2026-46333	49	Oracle Linux kernel (ptrace)
CVE-2026-46300	47	SUSE Linux kernel updates
CVE-2026-43500	38	90-day disclosure policy debate
CVE-2025-54518	30	SUSE Linux kernel updates
CVE-2026-33825	29	FortiGate VPN exploitation

Supply Chain

Supply-chain compromise was the strongest mover above the collection line, rising 85% as an attack type (369 to 684 mentions) and 194% as a technique (T1195, 157 to 462). One npm worm accounts for most of it.

Mini Shai-Hulud, the fourth-generation Shai-Hulud variant run by TeamPCP, ran through the quarter as a self-propagating npm worm that steals developer tokens and republishes itself across every package a compromised maintainer controls (Unit 42).

On 11 May it hit TanStack's build pipeline, publishing 84 malicious package versions in six minutes; on 12 May its operators released the source code, which spawned copycats; and on 19 May a single wave pushed malicious versions across 323 packages in Alibaba's AntV ecosystem, representing roughly 16 million weekly downloads, in a 22-minute burst (Snyk).

Two features drew most of the coverage: it was the first npm worm shown to ship packages carrying valid build-provenance attestations, and it hooked AI coding agents such as Claude Code for persistence. The activity surfaces in the malware table as Shai-Hulud (32 mentions), Mini Shai-Hulud (18, new) and the related Miasma payload (16, a June TeamPCP release).

Clop's Oracle E-Business Suite campaign is now a continuing-fallout story rather than a Q2 surge. Coverage of the Clop brand peaked at the Q4 2025 baseline (106 mentions) when the campaign broke, and by Q2 had fallen to 22.

The underlying facts stand: the actors exploited the zero-day CVE-2025-61882, extorting victims under the Clop brand, which researchers link to FIN11 though Mandiant cautioned it had not fully substantiated the actors' own claims. Confirmed victims include The Washington Post; 29 organisations had been named on the leak site by late 2025, with data leaked from 18 of them, and Halcyon put ransom demands as high as $50 million.

The pattern is the one Clop ran against MOVEit: compromise a widely deployed enterprise platform once, extort many downstream victims.

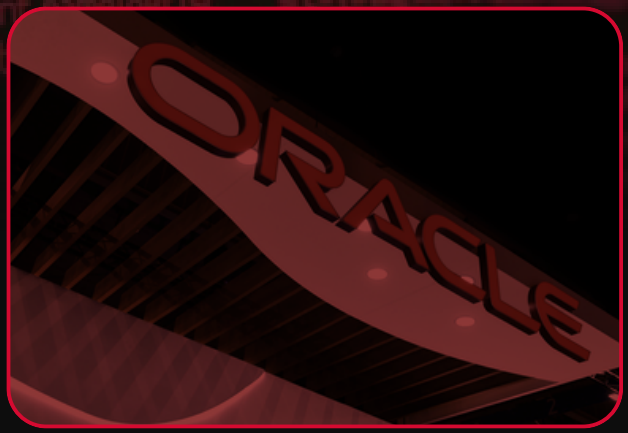

Both stories sit inside a wider acceleration. Sonatype counted 454,648 new malicious open-source packages in 2025, up 75% year on year, and the AI software supply chain has become its own strand, taking in the TeamPCP compromise of the LiteLLM AI gateway, critical flaws in the Langflow agent platform (CVE-2025-34291, added to CISA's KEV catalogue in May) and the PickleScan model scanner, and malicious models distributed through Hugging Face.

AI as weapon, target and stolen asset

The AI story ran through every part of the quarter and is treated in full in the companion brief; in summary, it ran on three tracks, and ThreatCluster's own data now measures each of them.

As an asset being stolen, the headline was Anthropic's 10 June letter to the US Senate Banking Committee accusing operators tied to Alibaba's Qwen lab of running 28.8 million Claude exchanges through roughly 25,000 fraudulent accounts between 22 April and 5 June to copy the model's capabilities, which Anthropic called the largest distillation campaign it has disclosed.

The figures are Anthropic's allegation, Alibaba denies them, and none is independently verified. Two days later the US Commerce Department restricted Anthropic's own Fable 5 and Mythos 5 models on national-security grounds.

As a weapon, open-weight models reached competitive parity on offensive-security work. Zhipu AI's GLM-5.2, released under an MIT licence on 13 June, beat Claude Code on Semgrep's IDOR benchmark (39% to 32% F1) at roughly a sixth of the cost, though Semgrep stressed the test is narrow.

OALABS documented the in-the-wild version: a low-skill operator in Addis Ababa used stolen Claude Code and Codex installs to breach at least 14 companies from prompts as thin as "recon this", with the agents raising only ten policy-violation blocks across more than 1,000 sessions. The Gentlemen's use of DeepSeek and Qwen to build its encryptor, noted above, is the same theme inside the ransomware economy.

As a target, the platform's data shows an AI-attack surface taking shape. Prompt injection rose 367% quarter on quarter as an attack type (6 to 28 mentions, climbing each month), and an AI-targeting malware category emerged around families such as Promptspy.

The research bore this out: the macOS.Gaslight prompt-injection implant aimed at AI malware-triage tools, and Mozilla's 0DIN proof-of-concept that walks Claude Code into opening a reverse shell from a repository containing no malicious code.

The consumer-facing cost is measurable: the FTC reported $3.5bn lost to imposter scams in 2025 inside a record $16bn of total fraud. The defensive side of the same technology also registered in the data: Anthropic's autonomous vulnerability-discovery programme, tracked in the platform as the Project Glasswing campaign, appeared new in Q2 at 28 mentions.

Malware and spyware Activity

The operational core is the commodity infostealer. StealC rose 173% and spiked in June (32 of its 41 mentions), Vidar rose 208%, Lumma and Lumma Stealer together reached 40, and Amadey arrived new at 25.

Family	Baseline	Q1	Q2	Q2 vs Q1
Pegasus	13	27	65	141%
StealC	9	15	41	173%
Vidar	5	13	40	208%
Cobalt Strike	11	10	32	220%
Shai-Hulud	25	10	32	220%
Amadey	0	0	25	New
Stuxnet	1	7	24	243%
Glassworm	17	16	23	44%
Graphite	3	2	22	Small Base
Lumma	6	5	20	300%
Lumma Stealer	26	11	20	82%
Mini Shai-Hulud	0	0	18	New

These feed the credential-and-session economy that initial-access brokers and the extortion groups above rely on, which is consistent with T1003 OS Credential Dumping (up 210%) and T1078 Valid Accounts (up 186%) ranking high in the technique data. Cobalt Strike remains the steady post-exploitation framework, joined by PlugX and ShadowPad on the espionage side.

Mercenary spyware led the table. Pegasus (NSO Group) rose 141% to 65 mentions and Graphite (Paragon) rose off a small base to 22, sustaining the attention these tools draw around journalist and civil-society targeting.

Stuxnet's presence at 24, front-loaded in April, is retrospective and analytical rather than operational, tracking the quarter's Iran and Israel military reporting; a high mention count does not imply live activity.

Among new families, the cleanest signals are the Shai-Hulud lineage (Mini Shai-Hulud, Miasma), SpectralViper (the OceanLotus, Vietnam-aligned backdoor first seen in June), and SprySOCKS.

An AI-targeting category is forming around families such as Promptspy. The broader shape is the long-tail expansion noted earlier: distinct malware families more than doubled since inception, so the named families here sit on top of a rapidly widening base rather than dominating it.

Sectors, geography and techniques

Government held the most coverage but rose in line with ingestion, so the real movers are financial (up 130%), energy (up 115% quarter on quarter and 422% above baseline), technology, education, transportation and defence, all clear of the collection line.

Healthcare and telecommunications rose below it, meaning their share of attention fell even as raw counts rose. Government, healthcare, technology, defence, insurance and critical infrastructure all accelerated month on month through Q2, which points to where the pressure was building into the summer.

Industry	Baseline	Q1	Q2	Q2 vs Q1
Government	463	670	959	43%
Financial	344	362	833	130%
Energy	121	294	632	115%
Healthcare	232	344	465	35%
Manufacturing	174	180	298	66%
Technology	115	142	276	94%
Education	84	116	234	102%
Telco	66	162	189	17%
Transportation	54	42	130	210%
Retail	141	78	116	49%
Defence	13	13	37	185%
Insurance	22	15	39	160%

Geography tracked live conflict. Russia, Ukraine and the United States led, followed by a Middle East cluster around the Iran and Israel military events (Iran 1,627, Israel 709, Lebanon 230, UAE 221, Turkey 197) and an India and Pakistan cluster (India up 199%, Pakistan up 184%) around Operation Sindoor.

The United Kingdom rose 158% to 415 and climbed each month. These count any tagging of a nation in an article and conflate attacker, victim and bystander.

Country	Baseline	Q1	Q2	Q2 vs Q1
Russia	317	656	2,260	245%
United States	298	649	1,812	179%
Ukraine	112	413	1,766	328%
Iran	109	889	1,627	83%
China	365	578	1,531	165%
India	149	238	711	199%
Israel	39	395	709	79%
North Korea	168	215	706	228%
Germany	107	236	557	136%
United Kingdom	48	161	415	158%

Phishing led by a wide margin; counting the parent T1566 with the spearphishing-link subtechnique puts phishing references above 1,700. The two exfiltration techniques together (T1041 and T1567, 1,284) reflect the extortion-first economy, where data theft is the product.

Supply Chain Compromise corroborates the 85% attack-type rise, and Valid Accounts and OS Credential Dumping near the top confirm that credential abuse, not novel exploitation, remains the common initial-access path for the most active groups.

Outside the top ten, web-shell activity (T1505.003) was the fastest riser, up 348%, tracking the edge-appliance exploitation above.

MITRE technique	Q1	Q2	Q2 vs Q1
Phishing	1,039	1,284	24%
Exfiltration Over C2 Channel	437	824	89%
Application Layer Protocol	257	493	92%
Command and Scripting Interpreter	219	474	116%
Supply Chain Compromise	157	462	194%
Exfiltration Over Web Service	205	460	124%
Spearphishing Link	261	453	74%
Valid Accounts	125	357	186%
OS Credential Dumping	105	326	210%
Exploitation for Privilege Escalation	140	316	126%

The campaign data ties the actor and geography threads together. Operation Epic Fury and Operation Midnight Hammer track the US and Iran military events, and Operation Sindoor the India and Pakistan cluster.

Contagious Interview is the continuing North Korean developer-targeting operation. Operation Endgame, the law-enforcement takedown line, surged late in the quarter (2 mentions in April, 3 in May, 28 in June).

Project Glasswing is Anthropic's autonomous vulnerability-discovery programme, and FortiBleed is the late-Q2 credential-exposure campaign against Fortinet firewalls covered above. Salt Typhoon and ClickFix both declined, consistent with their falling actor and technique coverage elsewhere in the report.

Campaign	Baseline	Q1	Q2	Q2 vs Q1
Operation Epic Fury	0	39	54	38%
Operation Sindoor	5	1	41	4000%
Operation Endgame	27	6	33	450%
Project Glasswing	0	0	28	new
Mini Shai-Hulud	0	0	20	new
Operation Ramz	0	0	18	new
ClickFix	18	17	13	-24%
Operation Midnight Hammer	0	2	12	500%
Contagious Interview	10	8	10	25%
Salt Typhoon	19	28	10	-64%

What to watch in Q3

Fragmentation is the structural condition, not a one-quarter blip: the affiliate churn that produced World Leaks and The Gentlemen is likely to keep producing short-lived extortion brands rather than a single dominant successor, and the ransomware-to-data-theft shift will keep pushing activity out of the ransomware tag and into breach and supply-chain coverage.

The Mini Shai-Hulud source code is public, so the technique of hijacking build pipelines to produce valid provenance is now available to any capable actor, and copycat waves should be expected.

Edge and management-plane appliances remain the most reliable mass-exploitation surface, whether through zero-days like the Cisco SD-WAN and PeopleSoft chains or credential-exposure campaigns like FortiBleed against Fortinet firewalls.

On the AI front, the gap between export-controlled Western models and freely downloadable open-weight models that match them on security tasks is the change defenders will be living with, and the ransomware economy's early adoption of DeepSeek and Qwen shows where that capability flows first.

The clearest momentum signals into Q3 are the ones still accelerating in June: StealC and the infostealer economy, web-shell and public-facing exploitation, data-breach coverage, The Gentlemen, and the Russian re-emergence of APT29 and Turla.

Sources

The source base widened considerably. Linuxsecurity led ingestion at 3,386 articles, ahead of Cybersecuritynews (1,303), Gbhackers (1,069), The Register (1,018) and The Guardian (897). The presence of YouTube (726) and Hacker News (661) high in the list is worth noting for anyone reading the entity counts as a proxy for ground truth, since both carry a different

signal-to-noise profile than vendor research or national CERTs.
No single outlet drives the numbers. The largest source supplied 3,386 of 85,433 articles, about 4% of intake, and the seven largest together about 11%.

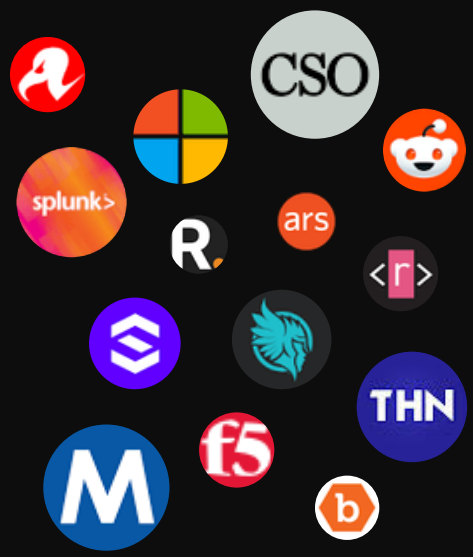

The remaining 10,000-plus sources averaged around seven articles each over the quarter, and that long tail is where the breadth findings come from: 1,174 distinct malware families and 583 nation-state groups surface because regional outlets, researcher blogs and niche feeds are collected alongside a mainstream press that concentrates on the head of the distribution.

Mention counts weight every article equally, so a syndicated feed item counts the same as a vendor deep-dive or a CERT advisory. Two features of the pipeline limit the distortion: clustering collapses duplicate coverage of the same story, which is why the cluster count rose only 9% against a 56% rise in articles, and alias merging keeps an entity's count from splitting across naming variants. What the counts cannot do is grade the quality of the reporting behind them, which is one more reason the report leans on concentration and trend figures, which are less sensitive to the source mix, rather than on any single quarter's raw totals.

Methodology and caveats

- Counts are distinct-article mentions, merged to canonical names across known aliases, and represent reporting attention rather than real-world incidence.
- Comparisons use three periods: a Q4 2025 baseline (the first full quarter of operation), Q1 2026 and Q2 2026. Intake rose 56% in articles and 57% in sources between Q1 and Q2, which mechanically lifts most raw counts; the report flags where a rise clears that line.
- Baseline comparisons for MITRE techniques and the highest-frequency countries over-read, because canonical tagging was still being rolled out in late 2025, so the baseline period carried few structured technique tags; for those series the quarter-on-quarter figure is the reliable trend.
- Percentages off very low counts are flagged as small-base.
- First-seen counts include legacy families surfacing under first-time tags; de-duplicated estimates are used for the genuinely-new findings (roughly 63 of 84 raw new ransomware entrants, and 377 of 415 new nation-state clusters, are genuine after de-duplication; the malware new-entity count is not de-duplicated and is treated as an upper bound). Where a specific operational fact is stated, it is attributed to a named primary source.
- Victim totals drawn from leak-site trackers are real postings, not confirmed totals of real-world incidence, and may include duplicate or unverified claims. Attribution statements reflect the confidence levels of the named researchers, not an independent ThreatCluster assessment.

ThreatCluster

As featured on

AP · International Business Times · abc NEWS · CBS · Noah News PREDICTIVE INTELLIGENCE

threatcluster.io
hello@threatcluster.io
71-75 Shelton Street, Covent Garden, London, WC2H 9JQ