



ARMATURE
SYSTEMS

APR — JUN 2026

Global Cyber Threat Landscape Report

Table of Contents

1	EXECUTIVE SUMMARY	2
2	THREAT ACTOR TRENDS	3
3	ARMATURE THREAT RESEARCH	5
4	ATTACK VECTORS, TECHNIQUES & EXPLOITATION TRENDS	8
5	EMERGING THREATS & TECHNOLOGY RISKS	10
6	RECOMMENDATIONS	13

1 Executive Summary

Why find a backdoor when you can walk through the front door? That was Q2 in a sentence. The quarter's most effective attacks didn't rely on novel malware. It relied on trusted vendor connections, valid credentials harvested from users, continued vishing campaigns, and internet facing device vulnerabilities.

Two Palo Alto PAN-OS firewall zero-days and two Fortinet FortiClient EMS remote code execution flaws were chief among the exploits in H1. The devices bought to enforce the perimeter became the way through it. In several cases the breach was never logged by the firewall, giving the attacker free rein to move deeper into the internal network before the SOC had any signal to act on.

4

Edge-appliance zero-days exploited across H1 2026

5

Major threat-actor groups profiled

15

Median MITRE ATT&CK techniques per AI-assisted intrusion

275M

Records ShinyHunters claims from the Instructure/Canvas breach

Against that backdrop, three themes shaped the quarter.

- **AI accelerated exploitation has moved to the network edge.** The window between vulnerability disclosure and active exploitation has compressed from months to hours, with generative AI now assisting adversaries across a median of fifteen distinct MITRE ATT&CK techniques per threat actor (Verizon 2026 DBIR, in partnership with Anthropic). The targets of choice were edge devices (PAN-OS, FortiClient EMS, Cisco Catalyst SD-WAN and Ivanti) because a single internet-reachable appliance offers root-level access, in the worst cases with no credentials required.
- **The trusted third party has become the largest blast radius.** Q2 was dominated by identity-led intrusions into SaaS platforms, most visibly Salesforce, and by extortion crews working in tandem, exemplified by the Qilin-then-ShinyHunters multi-group extortion of Sysco and the record-setting Instructure/Canvas breach affecting up to 275 million user records.

- **Ransomware is increasingly sector targeted, set against sustained geopolitical cyber operations.** Ransomware concentrated on operationally fragile sectors (manufacturing, pharmaceutical supply, and food and agriculture) even as the criminal ecosystem fragmented. Layered over this, the cyber dimension of the 2026 Iran war matured from acute disruption into persistent targeting of operational technology directed by Iran's IRGC and MOIS (Ministry of Intelligence and Security), alongside state-aligned ransomware collaboration

2 Threat Actor Trends

Q2 2026 was defined less by a single dominant actor than by the convergence of financially motivated extortion crews and state-aligned operators around the same techniques: identity abuse, edge-device exploitation and operational disruption. The most significant development was the rise of collaborative, multi-group extortion, in which a ransomware operator and a data-theft brand target the same victim in sequence.

1. ShinyHunters | Defining Extortion Brand of the Quarter

ShinyHunters was the most consequential financially motivated actor of Q2, responsible for the Instructure/Canvas, Carnival and Charter Communications intrusions and a co-claim against Sysco. Its method is consistent and effective: continued voice-phishing (vishing) and pretexting against helpdesk and employee accounts to compromise cloud identity, frequently Microsoft Entra, and pivot into connected SaaS, above all Salesforce. The group operates fluidly alongside Scattered Spider and Qilin affiliates under the loosely branded “Scattered Lapsus\$ Hunters” collective and was later observed exploiting Oracle PeopleSoft vulnerabilities to exfiltrate data and deploy ransom notes. We covered ShinyHunters activities as well in our Q1 report, and they will continue to be a threat in Q3.

2. Qilin (Agenda) | Most Prolific, Increasingly Challenged

Qilin — a Russian linked ransomware-as-a-service operation first identified in 2022, remained the most prolific group by observed victim volume. its open affiliate recruitment model produces high victim counts alongside notably high rates of non payment, leaving it prolific but comparatively unprofitable.

3. The Gentlemen | Fastest Rising Operator

First appearing in the second half of 2025, The Gentlemen became one of the most active groups, and sustained a high operational tempo through Q2 across manufacturing, professional services, education and logistics. The speed of that ascent strongly indicates experienced affiliates and operators behind the brand, a recurring pattern as personnel migrate between dissolving and emerging RaaS operations.

4. Akira & DragonForce | Maturing RaaS, Shifting Affiliates

Akira, one of the longest-running active RaaS groups (since 2023), and DragonForce continued to consolidate the post-RansomHub affiliate base under a franchise-style model, having absorbed operators following RansomHub's collapse in April 2025. The net effect is an ecosystem that is fragmenting rather than consolidating, with smaller, shorter-lived crews that are harder for law enforcement to track.

5. Iran-Nexus Actors | Sustained OT/ICS Targeting & State Ransomware Collaboration

Following the outbreak of the 2026 Iran war, the cyber dimension persisted through Q2. In late March, the IRGC-aligned CyberAv3ngers cluster (CL-STA-1128) shifted from internet-connected Unitronics PLCs to targeting Rockwell Automation OT environments. The MOIS-linked Handala persona, assessed as responsible for the destructive March wiper attack against the company Stryker, continued hack-and-leak operations despite an FBI domain seizure. Notably, the Iranian state-aligned ransomware group Pay2Key.I2P was observed offering profit shares to affiliates willing to attack U.S. and Israeli targets, a direct fusion of state objectives and criminal monetization that raises spillover risk for organizations with ties to the conflict.

Source: Recorded Future

3 Armature Threat Research - Identity Is The New Battleground


Attackers have shifted from breaking into systems to logging into them. Rather than exploit a technical flaw, they increasingly target the user's identity, using social engineering to impersonate a legitimate user and gain access that looks authorized. Their methods range from sophisticated to surprisingly simple, but share one trait: they exploit human trust rather than a technical vulnerability.

One such technique is device code phishing. What makes device-code phishing particularly noteworthy is its combination of social engineering with a living-off-the-land (LotL) approach. While LotL techniques have existed for many years, the emergence of attacks that blend LotL with social engineering presents a significant challenge for defenders. Because these attacks rely on legitimate services, trusted authentication mechanisms, and user interaction rather than malware, they are difficult for both signature-based and behavioral-based detection technologies to identify. Social-engineering attacks have always been challenging to detect, and combining them with LotL techniques further increases their effectiveness while reducing the likelihood of detection, even in environments protected by modern security technologies.

During our Managed Detection and Response (MDR) operations and threat-research activities, we identified multiple device code phishing attacks in which the threat actors demonstrated a strong emphasis on operational security. Furthermore, following our defensive actions, we observed that the attackers dismantled the limited infrastructure they had used. Armature threat research identified several IoCs related to this attacker. Before discussing the IoCs, let us first understand device-code phishing.

ATTACK TECHNIQUE · IDENTITY & ACCESS

Device Code Phishing



As diagrammed above in the image, following are the key steps executed to achieve a device code phishing attack.

1. Attacker requests a device code sent to the legitimate identity provider
2. Provider returns code and URL a short user code and verification link
3. Attacker delivers the code via lure email or chat impersonating IT support
4. Victim signs in and completes MFA on the genuine provider login page
5. Tokens issued to attacker session access and refresh tokens, not the victim
6. Attacker gains persistent access no password is ever captured

Because device-code phishing leverages legitimate Microsoft services, standard defensive solutions often fail to flag the activity as malicious. Furthermore, Armature research observed that email content-validation filters fail, since the included links point to legitimate URLs. The attack flow requires the target user to enter a code on an official Microsoft service, meaning even technically aware users frequently fall victim. Ultimately, traditional user-awareness training is largely ineffective against this specific type of attack.

FIELD	VALUE
IP address	43.157.2[.]149
Domain	qrtrac[.]link
URL 1	qrtrac[.]link/xkdH
URL 2	hxxps://docsxfiles.s3.us-east-2.amazonaws[.]com/Shared.Document.html

Mitigating device-code phishing starts with disabling the OAuth device-authorisation flow entirely for any user population or application that does not have a genuine need for it, since the flow is only required for input-constrained devices such as smart televisions, conference-room systems and certain IoT hardware. Where the flow must remain enabled, scope it tightly using Conditional Access policies that restrict device-code authentication to trusted networks, compliant devices or specific user groups, so that a code entered from an unexpected location or an unmanaged device is blocked outright. Strengthen user awareness as a parallel control, because the technique depends on a person being persuaded to type a code into a legitimate-looking prompt; staff should be trained to treat any unsolicited request to enter a device code as suspicious, especially when it arrives by email or chat and claims to come from IT support. On the detection side, monitor sign-in logs for device-code authentication events and alert on patterns that suggest abuse such as a code being requested from one location and redeemed from another, authentication from unfamiliar device types, or a sudden rise in device-code grants across the estate.

Finally, reduce the value of any stolen tokens by enforcing short token lifetimes, enabling continuous access evaluation so that sessions can be revoked quickly when risk is detected, and ensuring that compromised refresh tokens can be invalidated promptly once an incident is identified.

Source: Armature Threat Research

4 Attack Vectors, Techniques, & Exploitation Trends

The central analytic finding of Q2 is the inversion of the initial-access hierarchy. Vulnerability exploitation has displaced stolen credentials as the leading entry point, a strategic pivot toward systematic exploitation of unpatched, internet-facing systems at scale, accelerated by AI that collapses the time from disclosure to exploitation.

Most Significant Exploited Vulnerabilities

Edge and perimeter devices dominated the quarter’s exploitation, several attacked as zero-days before patches existed. The table below summarizes the most operationally significant CVEs added to CISA’s KEV catalog during Q2 2026.

CVE	PRODUCT / FLAW	Q2 2026 SIGNIFICANCE
CVE-2026-0300	PAN-OS User ID Auth Portal buffer overflow	Unauth root RCE; state-sponsored cluster CL-STA-1132 (tooling overlaps: Volt Typhoon, APT41).
CVE-2026-0257	PAN-OS GlobalProtect auth bypass	Cookie forged from public TLS cert.
CVE-2026-35616	Fortinet FortiClient EMS access control	Unauthenticated Remote Code Execution zero-day
CVE-2026-21643	Fortinet FortiClient EMS SQL injection	Earlier EMS RCE flaw (SQL injection) disclosed Q1; exploitation continued into Q2
CVE-2026-20122 CVE-2026-20128 CVE-2026-20133	Cisco Catalyst SD-WAN Manager	Vulnerabilities encompass improper privileged API usage, insecure credential storage, and unauthorized data exposure.
CVE-2026-31431	Linux kernel crypto (“Copy Fail”)	Local privilege escalation on multi-user/container hosts.
CVE-2026-34621	Adobe Acrobat & Reader prototype pollution	Actively exploited; client-side document attack surface.

Dominant Techniques (MITRE ATT&CK)

The techniques below were the most consistently observed across both nation state and criminal operations during the quarter.

Technique (ATT&CK ID)	Q2 2025 Observation
Exploit Public-Facing Application (T1190)	Leading initial access vector; centred on edge appliances, PAN-OS, FortiClient EMS, Cisco SD-WAN, Ivanti.
Phishing / Spear-Phishing (T1566)	AI enhanced lures at scale; AI-assisted initial access techniques are phishing related.
Phishing: Voice / Pretexting (T1566.004, T1598)	Vishing and pretexting resurgent; ShinyHunters/ Scattered Spider help desk compromise
Valid Accounts (T1078)	SaaS/SSO identity abuse.
Supply Chain Compromise (T1195)	Third-party involvement in breaches; shared vendor and SaaS integration blast radius.

Sources: CISA KEV catalog & Recorded Future.

5 Emerging Threats And Technology Risks

Agentic AI, Shadow AI Governance, And Vulnerability Pressures

Three AI-driven forces defined the emerging-threat picture this quarter. First, autonomous “agentic” AI is beginning to appear in attack chains, adapting tactics at machine speed and eroding the viability of human-speed response. Second, shadow AI has become a measurable data-leakage category, as employees route sensitive data through unapproved AI tools that the security team cannot see. Third, AI-assisted vulnerability discovery and exploitation are outpacing the ability of defenders to patch, a pressure examined in detail below.

Agentic AI is moving from demonstration to attack chain

Autonomous agents are beginning to execute multi-step intrusions with limited human direction: enumerating a target, selecting a technique, and adapting when a step fails, at a tempo no human operator can match. The significance for Q2 is not that agentic attacks are yet common, but that they compress the defender's reaction window toward zero. Controls that assume a human-paced adversary, such as manual triage queues and business-hours response, degrade against an attacker that does not tire and does not wait.

Shadow AI has become a measurable data-leakage category, not a hypothetical

The risk is no longer that employees might one day misuse AI; it is that they already route sensitive data through tools the security team cannot see. Staff reach AI services through personal accounts and unmanaged browser sessions, pasting source code, customer records, and internal documents into models that retain and train on the input. Because the traffic looks like ordinary web use and the accounts are not corporate, standard data-loss prevention and logging never fire. Shadow AI is the direct descendant of shadow IT, except that the data does not just leave the perimeter, it leaves in a form the model may reproduce for someone else. The governance gap is structural, because adoption inside the business is outpacing the security team's ability to inventory it, let alone guardrail it.

The pressure to patch faster is real, but speed without prioritization does not scale

Two forces are colliding. On the supply side, AI-assisted vulnerability discovery is flooding the queue: Microsoft alone disclosed a record 622 vulnerabilities in a single July 2026 Patch Tuesday, roughly triple the prior month's total. On the demand side, AI-assisted exploit development has collapsed the time from disclosure to working exploit from roughly 125 days in early 2025 to about half a day by April 2026. The result is a treadmill no organization wins by volume: 82 percent of organizations already carry vulnerabilities left open for more than a year. Patching everything faster fails as a strategy, because everything is growing faster than any team can remediate.

Sources: Recorded Future and Armature Threat Research

5 Emerging Threats And Technology Risks

Exploitability, not theoretical severity, is the prioritization signal that scales

A CVSS score describes how damaging a flaw could be in principle; it says nothing about whether anyone is actually using it. Two signals close that gap: EPSS, the Exploit Prediction Scoring System, which estimates the probability a vulnerability will be exploited in the near term, and CISA's Known Exploited Vulnerabilities catalog, which lists what is confirmed under active attack. Prioritizing by exploitability rather than raw severity is the difference between a program that reduces real risk and one that burns effort proving how much it patched. Even NIST has fallen behind: it enriched only about half of new CVEs between April and June 2026, so defenders can no longer assume an authoritative severity score even exists. That is exactly why exploitability signals like EPSS and CISA KEV, not CVSS alone, have to drive prioritization.

The AI threat surface is now wide enough to require its own map

Taken together, these trends do not describe a single new threat; they describe a new breadth of exposure. AI is at once a weapon pointed at the enterprise, an ungoverned channel out of it, a dependency embedded inside it, and a multiplier on the vulnerability backlog. The figure below maps the four fronts a defender must now cover at the same time. The strategic point is not any one box; it is that all four are expanding at once, and none is optional.

What Q2 2026 added to the defender's scope: four expanding fronts, none optional

AI ENABLED OFFENSE

- Agentic attack chains adapt at machine speed
- Exploit development: 125 days to under 1 day
- Automated vishing and deepfake pretexting
- AI assisted reconnaissance at scale

SHADOW AI

- Unapproved tools via non-corporate accounts
- Sensitive data pasted into public models
- No DLP coverage or audit visibility
- Guardrails lag behind adoption

AI SUPPLY CHAIN + MODEL RISK

- Third-party AI vendors, plugins, and agents
- Tool tokens and grants widen access
- Model and training-data integrity
- Every integration expands blast radius

VULNERABILITY PRESSURE

- 622 Microsoft CVEs in one month (Jul 2026)
- 82% carry vulnerabilities open over a year
- Prioritize by exploitability: EPSS + CISA KEV
- Disclosure-to-exploit window measured in hours


Supply Chain & SaaS Identity as Systemic Risk

The most consequential shift this year is the migration of systemic risk away from individual networks and towards the identity fabric that binds modern enterprises to their suppliers and SaaS platforms. For defenders, this reframes the priority from protecting the perimeter to continuously governing identity and third-party trust, treating every token, grant and integration as an asset that must be inventoried, scoped to least privilege, monitored for anomalous use and rapidly revocable when a supplier is breached.

Operational Technology Under Geopolitical Pressure

The convergence of geopolitical conflict and OT targeting is the quarter's most strategically concerning trend for industrial and manufacturing clients. IRGC-aligned actors expanded from internet-connected PLCs to Rockwell Automation environments, and the willingness of state-aligned groups to incentivise criminal attacks on Western targets removes a layer of restraint.

Sources: Recorded Future and Armature Threat Research



6 Recommendations

Critical Priority

Treat edge and security appliances as a crown-jewel attack surface

Inventory every internet-facing firewall, VPN, and management server. Restrict administrative and User-ID/Captive portals to trusted zones.

Harden SaaS and cloud identity against vishing-led intrusion

Enforce phishing-resistant MFA across all identity surfaces. Eliminate authentication override, cookie, and certificate reuse configurations.

Operationalize the infostealer-to-ransomware early warning signal

Treat infostealer infections and credential leak events as ransomware precursors. Where exposure is detected, rotate all affected identity artefacts: session tokens, API keys, and service account credentials.

Close the device-code and OAuth consent phishing path

Disable the OAuth 2.0 device-authorization grant flow where it is not required. Where it is required, scope it with Conditional Access, enforce short access-token lifetimes, and enable Continuous Access Evaluation so stolen tokens are revoked in near real time.

High Priority

Segment and govern OT/ICS environments

Iran-nexus targeting of Rockwell and PLC infrastructure is documented. Isolate operational technology from IT, monitor internet-facing OT, and remove default credential and exposed portal conditions.

Establish an AI security governance framework

Inventory sanctioned and shadow AI use. Apply data loss prevention controls to generative AI egress.

Formalize third-party and supply-chain risk management

Impose contractual breach notification clauses and least-access requirements on vendors. Require evidence of MFA remediation, and adopt SBOM practices for software dependencies.

Accelerate vulnerability remediation cadence

The disclosure-to-exploitation window is now measured in hours. Move from periodic patch cycles to continuous, risk-weighted remediation, prioritizing internet-facing assets, and sequencing by active-exploitation signal (EPSS and CISA KEV) rather than CVSS alone.

Prepare defenses for machine-speed, AI-driven attacks

Automate detection and containment for the controls most exposed to a shrinking reaction window. Add continuous monitoring with autonomous or semi-autonomous response for identity, edge, and endpoint, so containment does not depend on human reaction time



ARMATURE
SYSTEMS

AUTHORS



Jared Mendenhall
CISO - Armature Systems



Nadaraja Himasuthan
Threat Researcher- Armature Systems

ABOUT ARMATURE SYSTEMS

Armature Systems is a cybersecurity systems integrator that helps organizations design, build, and operate network and security programs across the full lifecycle.

We work hands-on alongside internal teams to deploy the right technology, build effective processes, and deliver measurable security outcomes across network, cloud, identity, compliance, and threat response.

Ready to modernize your security program? Visit armaturesystems.com