PANOP

PANOP RESEARCH · THREAT BRIEFING

Threat Actors 2026:
Who Is Actually Attacking You

A briefing on the four threat actor categories and where their boundaries have stopped holding: why the same group carries four different names, how far attacker tempo has collapsed, what the Swiss mandatory reporting data shows about who actually turns up, and the ninety days of work that follows from a twenty-nine minute breakout time.

EDITION SEPTEMBER 2026

PANOP RESEARCH · PANOP.IO

PANOP

01 · EXECUTIVE SUMMARY

Attribution matters less than tempo

Most threat actor writing is either a taxonomy with no numbers or a list of group names with no structure. The useful finding in the 2026 data is neither: it is that the interval between an attacker getting in and moving deeper has collapsed to well under an hour, while the intrusions that persist longest sit on the devices least likely to be generating logs.

THE POSITION IN ONE PARAGRAPH

You are unlikely to know who is in your network while that knowledge would still change your response. Average breakout time is now twenty-nine minutes, access is handed between criminal groups in a median of twenty-two seconds, and the categories that once implied a toolset no longer do — state actors deploy commodity malware and criminals borrow state tradecraft. The leverage has moved from responding faster to exposing less.

29 min

Average eCrime breakout time, down from 48 minutes in 2024

CrowdStrike 2026 Global Threat Report

22 sec

Median hand-off from access broker to the follow-on group

Mandiant M-Trends 2026

122 days

Median dwell for espionage and DPRK IT worker intrusions

Mandiant M-Trends 2026

32%

Of intrusions began with exploitation of an internet-facing system

Mandiant M-Trends 2026

THE THREE FINDINGS THAT SHOULD CHANGE A PLAN

Categories converged

Tooling no longer identifies the adversary. ENISA documents state-nexus groups running ransomware and commodity infostealers, and criminal groups adopting state-style social engineering. **Do not size a response to a guess about who it is.**

Dwell time reversed for a reason

Global median dwell rose to **14 days from 11**. Internal detection actually improved. The average worsened because long-dwell espionage on edge devices grew, and most programmes are not instrumented for it.

WHAT LEADERSHIP HAS TO DECIDE

01 **Do we know what is reachable from the internet today?**
Edge devices are simultaneously the leading exploitation target and the favoured hiding place for long-dwell intrusions.

02 **Do those devices produce logs we collect?**
VPNs, firewalls and boundary appliances typically lack standard telemetry, which is precisely why dwell on them is measured in months.

03 **What happens when someone calls the help desk?**
Voice phishing is now the second most common initial vector, at 11%, while email phishing fell to 6%. The target is a workflow, not an inbox.

04 **Does our response plan assume an obsolete clock?**
Any process written against an hours-long intrusion window is planning for an adversary that no longer exists.

The gap. Detection and response remain necessary and are no longer sufficient on their own. At twenty-nine minutes, the decisive work happens before the intrusion — in what is exposed, what is reachable and what is actually exploitable.

SOURCES CrowdStrike, 2026 Global Threat Report, 24 February 2026. Mandiant and Google Cloud, M-Trends 2026. ENISA Threat Landscape. NCSC semi-annual report 2026/I. Full references on page 9.

02 · THE NAMING PROBLEM

Why the same group carries four different names

This has to be dealt with before any of the numbers make sense, because it is the most common reason threat intelligence is misread inside an organisation.

FOUR TAXONOMIES, ONE ADVERSARY

VENDOR	CONVENTION	NAME FOR ONE RUSSIAN STATE-NEXUS GROUP
Microsoft	Weather. Blizzard for Russia, Typhoon for China, Sandstorm for Iran, Sleet for North Korea, Tempest for financially motivated actors.	**Midnight Blizzard**
CrowdStrike	Animals. Bear for Russia, Panda for China, Kitten for Iran, Chollima for North Korea, Spider for criminal groups.	**COZY BEAR**
Mandiant	APT numbers for graduated groups, UNC numbers for uncategorised activity clusters.	**APT29, UNC2452**
Unit 42	Constellations. Applied across both state and criminal actors.	—

The criminal group most people call Scattered Spider is Octo Tempest to Microsoft and Muddled Libra to Unit 42. In June 2025 Microsoft and CrowdStrike published a joint mapping, later joined by Mandiant and Unit 42, deconflicting more than **80 adversaries**. It is deliberately not a single naming standard — each vendor keeps its own system and its own telemetry — but it is a usable translation table, published as JSON in Microsoft's MSTIC repository.

TWO CONSEQUENCES WORTH ACTING ON

Do not count groups across feeds

Asking how many distinct groups attacked you, using two vendor feeds without deconfliction, will overstate the answer. The mapping exists precisely so that this work is not repeated by every practitioner.

A name is a behaviour label

A group name is a cluster of observed activity, not a legal identification of an organisation. Vendors also differ on where one cluster ends and another begins. Treat the name as a description of tradecraft and you will use it correctly.

WHAT THE NAME DOES NOT TELL YOU

Attribution is slow, contested and arrives after the decisions that matter. It is valuable for policy, law enforcement and long-range strategy. It is close to useless as an input to containment, because the intrusion will have progressed well before any confident label is available — and, as the next section shows, the label would no longer reliably predict the tooling in any case.

Practical rule. Use actor names to communicate tradecraft and to correlate reporting. Do not use them to triage, to size a response, or to decide whether an alert deserves attention.

SOURCES Microsoft Security, Announcing a new strategic collaboration to bring clarity to threat actor naming, 2 June 2025. CrowdStrike press release, 2 June 2025. Microsoft MSTIC public threat actor naming feed.

PANOP

03 · THE FOUR CATEGORIES

Motivations still hold. Toolsets no longer do.

The categories remain a good description of why an adversary is present. What has broken down is the assumption that each owns a distinct toolset — the finding ENISA calls convergence.

THE FOUR CATEGORIES

CATEGORY	WHAT SUCCESS LOOKS LIKE TO THEM	TYPICAL DWELL
State-nexus espionage	Intelligence collection, and above all remaining undiscovered. Increasingly persistent on edge devices that lack standard telemetry.	**122 days**
Cybercriminal	Fast monetisation through encryption, extortion or fraud. Speed is the design goal, and the ecosystem is specialised into access and follow-on roles.	**Minutes to days**
Hacktivist	Attention. Visible disruption that can be screenshotted and claimed, which makes volume a poor proxy for impact.	**Transient**
Fraudulent worker	Being paid a salary while collecting access and data. North Korean IT worker operations are the documented case at scale.	**122 days**

CONVERGENCE, IN BOTH DIRECTIONS

State actors using criminal tooling

ENISA documents DPRK-nexus **Kimsuky** using the ClickFix technique, **Andariel** linked to Play ransomware as an apparent affiliate or access broker, **Moonstone Sleet** leveraging Qilin, and China-nexus **Mustang Panda** using RA ransomware. **APT29** and **Sandworm** have been observed on commercial residential proxy networks, sharing hosting with criminals and deploying commodity infostealers.

ENISA Threat Landscape

Criminals using state tradecraft

The traffic runs the other way too. The financially motivated group **FIN6** adopted fabricated LinkedIn personas and job-application lures — a playbook borrowed directly from North Korean operations, and now a standard technique for reaching developers.

ENISA Threat Landscape

The supply chain, from both sides

The Swiss NCSC makes the same point about open-source package manipulation, noting that documented cases span financially motivated groups such as **TeamPCP** and state-sponsored actors including **Lazarus Group** and **APT29**. One technique, two categories.

NCSC semi-annual report 2026/I

Why the blend is rational

Commodity tooling is cheap, effective and provides deniability, because its presence implies a commodity adversary. Criminal infrastructure is rented rather than built. Convergence is not sloppiness on the adversary's part; it is a sensible economy that happens to degrade your inference.

What this means for defence. You cannot infer an actor's category from their tooling. A commodity infostealer is no longer evidence of a commodity adversary, and a ransomware note is no longer evidence that money was the objective.

SOURCES ENISA Threat Landscape, on convergence and cross-over between state-nexus and cybercriminal tradecraft. NCSC semi-annual report 2026/I, on open-source package manipulation. Mandiant M-Trends 2026, on dwell time by intrusion category.

PANOP

04 · TEMPO

The measurement that should change your plans

This is where the 2026 data differs materially from previous years. Every figure below moved in the same direction, and they compound.

SPEED, MEASURED

MEASUREMENT	2026 FIGURE	PRIOR
Average eCrime breakout time, initial access to lateral movement	**29 minutes**	**48 min in 2024**
Fastest breakout on record	**27 seconds**	—
Median hand-off from access broker to follow-on group	**22 seconds**	**Over 8 hours in 2022**
Investigations showing the access hand-off model	**9%**	**4% in 2022**
Detections involving no malware at all	**82%**	—
Increase in AI-enabled adversary activity	**89%**	**Year on year**

At twenty-nine minutes on average, an alert waiting in a tier-one queue has already been overtaken. In one intrusion observed by CrowdStrike, **exfiltration began four minutes after initial access**. The hand-off figure deserves separate attention: it no longer routes through a forum sale, so the delay that used to give defenders a window has simply gone.

DWELL TIME WENT THE WRONG WAY, FOR AN INSTRUCTIVE REASON

HEADLINE

14 days

Global median dwell time, up from 11 days in 2024. Read alone, this looks like defensive failure.

→

BUT

9 days

Internally detected intrusions improved, from 10 days. Defenders got faster at the cases they were built for.

→

WHILE

25 days

Externally notified cases more than doubled, from 11 days. Someone else told them.

→

BECAUSE

122 days

Espionage and DPRK IT worker cases, some beyond a year. This category grew.

The mix changed rather than the competence. Long-dwell espionage increasingly persists on **edge devices — VPNs, firewalls and boundary appliances that lack standard telemetry**. The global average got worse because a category most programmes are not instrumented for got bigger, which is a very different problem from defenders getting slower, and it has a very different remedy.

> **The uncomfortable arithmetic.** Twenty-nine minutes to lateral movement, twenty-two seconds to hand-off, and 82% of detections with no malware to signature on. Controls that depend on recognising something known, in time, are being asked to do something the clock no longer permits.

SOURCES CrowdStrike, 2026 Global Threat Report, 24 February 2026. Mandiant and Google Cloud, M-Trends 2026, drawing on over 500,000 hours of incident response in 2025.

PANOP

05 · HOW THEY GET IN

Exploitation leads. The phone is now second.

For the sixth consecutive year, exploitation of internet-facing systems was the leading initial infection vector in Mandiant's investigations.

INITIAL INFECTION VECTOR, WHERE ONE COULD BE IDENTIFIED

VECTOR	2025	2024
Exploitation of internet-facing systems	**32%**	**Leading**
Voice phishing	**11%**	—
Email phishing	**6%**	**14%**

Two shifts matter. Email phishing **more than halved**, while voice phishing rose to second place. Attackers moved to interactive, rapport-building social engineering aimed at help desks and IT support, because a conversation is far more resilient to automated technical controls than a message that can be filtered. Both sit under social engineering, but they require entirely different detection strategies, and only one of them is addressed by an email gateway.

THE EDGE IS THE TARGET

Where the exploitation lands

CrowdStrike found **40% of vulnerabilities exploited by China-nexus actors targeted edge devices**. The Swiss NCSC independently describes inadequately protected edge devices as a key gateway into affected networks, and as raw material absorbed into wider attack infrastructure.

Why they are attractive

Edge devices are internet-exposed by definition, frequently fall outside the change-control discipline applied to core business applications, and rarely produce the telemetry that would reveal an intruder. Exposure, weak governance and blindness in a single asset class.

A CASE WORTH READING IN FULL: STATIC TUNDRA

The NCSC report documents a Russian state-nexus actor reaching operational technology through unpatched end-of-life network devices. The failures are ordinary rather than exotic, which is what makes the case useful: remote access points were **exposed directly to the internet without multi-factor authentication**, the same passwords were reused across multiple devices, and in several cases equipment still carried **factory default credentials**. Firewalls, servers and OT appliances including controllers, remote terminal units, protection relays and human-machine interfaces were affected, some granting root privileges.

Having gained access, the attackers corrupted control device firmware so equipment entered an endless restart loop, deleted operating system files so devices could no longer boot, reset components to factory settings, and deployed wiper malware on Windows-based control computers. Unlike ransomware, a wiper is not a negotiating position.

> **The through line.** The leading way in is an exposed system with a known weakness, and the most damaging cases in this year's reporting turned on credentials and MFA on internet-facing devices — not on novel malware.

SOURCES Mandiant and Google Cloud, M-Trends 2026, on initial infection vectors. CrowdStrike, 2026 Global Threat Report, on edge device targeting. NCSC semi-annual report 2026/I, on edge devices, helpdesk impersonation and the Static Tundra case.

PANOP

06 · THE SWISS PICTURE

Flat in volume, fragmenting in structure

Switzerland has an unusually good public dataset since mandatory reporting for critical infrastructure began on 1 April 2025. In the first half of 2026 the NCSC received **27,128 voluntary reports**, down from 35,727 a year earlier, and **200 mandatory notifications**.

THE 200 MANDATORY REPORTS, BY ATTACK TYPE

REPORTED ATTACK TYPE	SHARE
Unauthorised access to systems (hacking)	**~26%**
Theft of login credentials	**13.5%**
Distributed denial of service	**12.7%**
Data leaks	**12.7%**
Ransomware and extortion	**~8%**

Public administration (19.4%) and IT and telecommunications (18.6%) reported the most. Two things are worth noticing. Ransomware's share is modest against the attention it receives. And the largest category, unauthorised access, largely means **compromised email accounts** subsequently used to run further phishing and fraud — an outcome of credential theft rather than a separate phenomenon.

RANSOMWARE: THE COUNT HELD, THE ECOSYSTEM SPLIT

79

Ransomware incidents, identical to the previous half-year

NCSC semi-annual report 2026/I

21 to 29

Active ransomware families operating against Swiss targets

NCSC semi-annual report 2026/I

35% to 19%

Akira's share, still the most documented group in the country

NCSC semi-annual report 2026/I

A stable incident count concealing a fragmenting supplier base is a harder intelligence problem than a rising one. Twenty-nine families produce more distinct tradecraft to track than twenty-one, with no change in the volume of work arriving.

TWO FINDINGS TO CARRY INTO A RISK ASSESSMENT

01 **No targeted sabotage against Swiss critical infrastructure, so far**

The NCSC has observed none to date, and is explicit that this depends on geopolitics. It notes Swiss infrastructure could become a proxy target because of its interconnection with neighbouring states, and points to the sabotage of distributed energy generation in Poland as the near case.

02 **Some victim listings are simply false**

Several Swiss organisations named on data leak sites during the period turned out not to be victims. Treat a leak-site listing as a claim to verify, with a named owner and a defined timeline, rather than as an incident.

Reading the Swiss data honestly. Voluntary reports fell, which is a measure of reporting behaviour rather than of threat. The mandatory notifications from critical infrastructure are the sounder signal, and they point at access and credentials rather than at exotic capability.

SOURCES NCSC / BACS, semi-annual report 2026/I, January to June 2026. Mandatory reporting for critical infrastructure operators in force since 1 April 2025.

PANOP

07 · LAW ENFORCEMENT

Pressure works, and it fragments the market

2026 saw sustained action against shared criminal infrastructure rather than individual gangs — the anonymity layer and the cash-out layer.

Operation Saffron — 19–20 May 2026

Dismantled First VPN, described by Europol as the most widely used anonymity service in the cybercrime underground and present in almost every major investigation it had supported. Led by France and the Netherlands across 27 countries, it seized 33 servers and arrested the administrator. The decisive part was recovering the **user database**: 506 users identified, feeding 21 other investigations.

AudiA6 takedown — 10 June 2026

Struck the cash-out layer instead, dismantling a laundering service that had processed more than **€336 million** between 2022 and 2025 and was linked to more than 15 separate ransomware investigations. Two administrators were arrested in Georgia.

This pressure has a second-order effect visible in the Swiss numbers. When affiliates are displaced — the NCSC cites the leak of the group The Gentlemen's own data in early May — they move to other groups rather than leaving. **Enforcement success shows up as fragmentation, not as fewer attacks.** It is a real gain, and it makes the tracking problem harder rather than easier.

08 · THE NEXT NINETY DAYS

You cannot out-respond a 29-minute breakout

Days 0–30 — Reduce what is reachable

Enumerate internet-facing systems and treat VPNs, firewalls and boundary appliances as the priority rather than the afterthought: they are at once the leading exploitation target and the favoured hiding place for long-dwell intrusions. Verify MFA on every remote access path, hunt reused and default credentials on network and OT equipment, and confirm end-of-life devices are genuinely gone rather than merely unsupported.

Days 31–60 — Instrument the gap

Establish what logging edge devices can actually produce and where it goes, because the absence of that telemetry is what turns an intrusion into a 122-day intrusion. Separately, rehearse the help desk: define an out-of-band verification step for password and MFA resets, since voice phishing is now the second most common way in and it targets that workflow specifically.

Days 61–90 — Rebase on tempo

Re-examine any response process whose assumptions predate a 29-minute breakout time, and decide in advance what gets contained automatically rather than after triage. Deconflict threat intelligence feeds against the joint vendor mapping so adversaries are not double-counted, and retire any control whose value depends on identifying the attacker early.

FIVE QUESTIONS FOR THE NEXT BOARD MEETING

01 Which internet-facing devices could an attacker reach today, and when did we last verify that list rather than consult it?

02 Do our edge appliances produce logs we actually collect — and if not, how would we ever detect a 122-day intrusion?

03 What happens when someone calls our help desk claiming to be an executive locked out of their account?

04 If a leak site named us tomorrow, who confirms whether it is true, and how long does that take?

05 How long does our response plan assume we have between initial access and lateral movement, and is that number still twenty-nine minutes or better?

The common thread. Breakout time, hand-off time and edge-device dwell time describe one problem from three angles: the attacker's speed is now a property of the ecosystem, while your reachable surface remains one of the few things still under your control.

PANOP

09 · REFERENCES AND METHOD

Every claim traces back to a public source

Use this page to check anything in the briefing, and to hand colleagues the primary material rather than the summary.

INCIDENT RESPONSE AND TELEMETRY

Mandiant and Google Cloud, M-Trends 2026, executive edition, drawing on over 500,000 hours of incident response conducted in 2025.

CrowdStrike, 2026 Global Threat Report, 24 February 2026, and its executive summary.

EUROPEAN AND SWISS AUTHORITIES

ENISA, Threat Landscape, on convergence between state-nexus and cybercriminal tradecraft and on sectoral targeting in the EU.

NCSC / BACS, semi-annual report 2026/I, January to June 2026, including the Static Tundra case study and Swiss ransomware statistics.

Europol, European Cybercrime Centre, on Operation Saffron and the AudiA6 takedown.

ATTRIBUTION AND NAMING

Microsoft Security, Announcing a new strategic collaboration to bring clarity to threat actor naming, 2 June 2025.

Microsoft MSTIC, public threat actor naming feed, mapping vendor aliases in JSON.

METHOD AND LIMITATIONS

This is a threat briefing, not legal advice, an attribution assessment or an assurance opinion.

Threat intelligence figures are vendor telemetry, not census data. Different vendors observe different customer populations, so absolute numbers should be read as direction and magnitude rather than precision.

Where two sources measure the same phenomenon differently, both are named rather than averaged. Figures are reproduced as published and were current at the edition date.

Actor names are clusters of observed activity, not legal identifications of organisations, and vendors differ on cluster boundaries.

Swiss voluntary report counts measure reporting behaviour as much as threat activity, and the NCSC notes the true incident count is likely higher than the reported one.

USING THIS BRIEFING

Circulate section 01 to risk, legal and board readers; sections 02 to 06 to security operations and architecture owners.

The ninety-day sequence in section 08 is written to be lifted directly into a workplan.

A web version with every source hyperlinked is published at panop.io/blog/threat-actors-2026.

Prepared by Panop Research, edition september 2026. This document is intended for awareness and briefing use. It is not legal advice and does not constitute an assurance opinion.