# The APT Paradox: Sophisticated Simplicity in Nation-State Cyber Operations (2024–2025), Trends, Detection Provenance, and Practical Gaps

**Raymond Andre Hagen**
Norwegian University of Science and Technology (NTNU), Norway

raymohag@stud.ntnu.no

**Abstract**: Advanced Persistent Threats (APTs) present a paradox in cybersecurity: sophisticated state actors use both zero day exploits and old social engineering tricks, maintaining complex infrastructure while exploiting basic misconfigurations. This study analyzes 60 verified APT campaigns from January 2024 to July 2025, providing an empirical snapshot of current threat actor behaviour, targeting patterns, and detection dynamics. Using a reproducible methodology with clear inclusion criteria based on state backing, persistence, and sophistica- tion indicators, we address four research questions: which actors are active (RQ1), what sectors they target and how this varies by actor (RQ2), which initial access methods dominate (RQ3), and who detects campaigns with what implications for visibility (RQ4). All data are archived in a public repository to enable validation and extension. Our findings reveal concentration among four primary state clusters: Russia (17 campaigns), China (16), North Korea (15), and Iran (9), accounting for 95% of attributed activity. Actor sector relationships show clear patterns: Chinese actors focus on telecommunications and government networks, Russians target diplo- matic infrastructure, North Koreans emphasize financial and cryptocurrency platforms, while Iranian operations cluster around regional events. Social engineering dominates initial access (40%), followed by web/network exploitation (21.7%) and N day exploitation (13.3%), with zero days appearing in only 8.3% of campaigns, chal- lenging assumptions about APT sophistication. Critical to defensive planning, we identify systematic detection gaps from vendor centric discovery that creates predictable blind spots in regions with limited commercial se- curity deployment and sectors using legacy infrastructure. The 18 month persistence of specific actor sector relationships indicates sustained rather than episodic interest, requiring continuous defensive evolution rather than one time responses. These findings require rethinking defensive strategies from isolated organisational responses to collaborative ecosystem approaches. The paradoxical nature of APT operations, advanced yet ba- sic, strategic yet opportunistic, reflects fundamental asymmetries in cyber conflict where attackers need only single successes while defenders must maintain continuous vigilance across expanding attack surfaces. Effective defense requires not just technical controls but coordinated, cross sector frameworks based on observed rather than theoretical threat behaviours. Scope and limitations: Findings reflect publicly reported activity within January 2024–July 2025 and may under-represent restricted disclosures (e.g., Five Eyes and allied operations). We analyse observable evidence in this window, not an exhaustive census. Practical implication: Although actors often chain techniques [Cybersecurity and Infrastructure Security Agency, 2025], we treat the first successful foothold as the decision-relevant initial access because it drives earliest containment and triage; later steps refine rather than replace these priorities.

**Keywords**: Advanced persistent threat (APT), Cyber threat intelligence, Initial access, Threat actor, Detection provenance

## 1. Introduction

Advanced Persistent Threats (APTs) are marked by strategic intent, multiphase operations, and sustained access against defined targets. Typically linked to nation state operators or comparably resourced proxies, such campaigns pursue long term collection, prepositioning, or disruption. Understanding who is active, what they target, how they gain initial access, and who brings activity to light is central to research and defense. The paradoxical nature of modern APT operations, sophisticated in capability yet often simple in execution, challenges traditional threat models and defensive assumptions. From January 2024 to July 2025, publicly reported APT activity remained high while actors retooled tradecraft. Although incidents receive extensive coverage through mostly vendor driven reporting, reproducible measurement over a clearly bounded period remains scarce, creating systematic visibility gaps that shape our understanding of the threat landscape. We address this by analyzing a verified corpus of campaigns that meet the APT definition Hagen and Helkala [2024], assembled from open sources and archived in a blinded OSF repository [Blinded, 2025]. All observed campaigns in the window are archived for transparency, while primary statistics use only the analytical subset.

Our analysis is organized around four questions: who is active (RQ1), what is targeted and how targeting varies by actor (RQ2), which initial access methods and tools dominate (RQ3), and who detects and reports campaigns, and what systematic gaps this creates (RQ4). Our analysis reveals that APT actors operate across a spectrum of sophistication, that detection ecosystems create predictable blind spots, and that effective defense requires collaborative approaches that go beyond organisational boundaries.

**Definition 1** (Advanced Persistent Threat (APT))**.** Following [Hagen and Helkala, 2024], we use *APT* to denote actors with (i) state backing or state level capability, (ii) sustained, campaign level operations against specific targets over time, and (iii) advanced tradecraft (e.g., bespoke malware, identity centric operations, supply chain compromise, or exploitation of public facing applications). This definition guides inclusion and coding in §3.2.

**Scope & limitations (summary).** Our dataset covers publicly reported APT campaigns during January 2024–July 2025 and may under-represent restricted disclosures from intelligence alliances. We analyse observable initial access and follow-on steps, not an exhaustive census of all operations. These constraints can skew prevalence estimates and regional coverage; we discuss implications and mitigation in §3.4.

## 2. Background and Related Work

Advanced Persistent Threats (APTs) are commonly distinguished by stealth, persistence, and strategic targeting. Unlike opportunistic or purely financially motivated intrusions, APT campaigns are typically linked to nation states or highly capable groups operating with geopolitical objectives [Chen et al., 2014, Rid, 2013, Valeriano et al., 2018]. Such operations seek long term access, sensitive data exfiltration, and, in some cases, disruption of critical functions. Two widely used lenses structure how practitioners describe and defend against APTs. The Cyber Kill Chain [Hutchins et al., 2011] models intrusion progress across phases and highlights opportunities for early detection, while MITRE ATT&CK [Strom et al., 2020] taxonomizes adversarial techniques and supports mapping between observed behaviour and defensive controls. These frameworks enable consistent language and cyber threat intelligence (CTI) sharing, but they are not a substitute for *empirical measurement* of which actors are active, which sectors are targeted, which techniques dominate, and who detects activity at a given time. Academic work has illuminated the geopolitical logic of state cyber operations. Valeriano and Maness [Valeriano et al., 2018] characterize cyber conflict as bounded by restraint and signaling, and Rid [Rid, 2013] situates operations within broader statecraft. Much of the empirical literature, however, remains retrospective and case centric, for example Stuxnet, APT1, and SolarWinds [Langner, 2011, Willett, 2021], with fewer recent, reproducible datasets that compare actors, sectors, and techniques within the same period.

**Present contribution.** We assemble a verified corpus of publicly reported and attributed campaigns observed between January 2024 and July 2025 and analyze the 60 records that meet our APT definition. All observed campaigns in the window are archived for transparency in a blinded OSF repository [Blinded, 2025], while borderline cases are retained for sensitivity checks but excluded from headline statistics. Each record is coded for actor attribution, targeted sectors and regions, initial access or TTPs, and detection source. Our analyses use descriptive statistics and cross tabulations aligned with four questions: who is active (RQ1), what is targeted and how targeting varies by actor (RQ2), which methods and tools dominate (RQ3), and who detects and reports campaigns (RQ4). Documentation is retained in the online archive to support review and replication. Compared to prior case oriented studies, this paper provides a recent, cross actor and cross sector snapshot of global APT activity over an 18 month window, grounded in observable evidence and designed for reproducibility.

**Conventional APT wisdom and uniform-treatment assumptions.** Operational guidance and standards often prescribe uniform first-line controls that fit most organisations and sectors. Agency advisories and widely used frameworks reinforce this baseline because they must be broadly applicable National Institute of Standards and Technology [2024], Cybersecurity and Infrastructure Security Agency [2023], iso [2023], European Union Agency for Cybersecurity (ENISA) [2025]. Our evidence indicates better early outcomes when priorities are calibrated to observed actor behaviour and sector exposure: where social engineering dominates, emphasise identity-centric containment and user-facing controls; where n-day reuse is common, accelerate patch triage and change management; where living-off-the-land is typical, bias monitoring toward abuse of native tools rather than malware signatures. [European Union Agency for Cybersecurity (ENISA), 2025] We therefore retain the baseline, but argue for a thin, evidence-led overlay tuned to actual APT economy-of-effort patterns observed in our window.

## 3. Methodology

### 3.1 Dataset Rationale and Scope

We assemble a verified corpus of **74** publicly reported campaigns for replication. An explicit inclusion gate yields the analytical subset of **60** campaigns that meet our APT definition, used for all statistics; the remaining **14** records are retained for sensitivity checks but excluded from headline results. Campaigns are included if they were ongoing on 1 January 2024 or initiated between 1 January 2024 and 1 July 2025. The dataset is frozen on 25 August 2025.

### 3.2 Data Availability

The dataset and documentation are available in a blinded OSF repository for peer review [Blinded, 2025].

**Inclusion criteria and attribution handling.** All observed campaigns within the window are archived for transparency [Blinded, 2025], while primary statistics use only the 60 records that satisfy our APT definition. State proxies count as APTs only when independent reporting shows tasking, command, funding, or operational control by a state service; ideological alignment alone is insufficient. We record attribution confidence to separate epistemic uncertainty from non APT activity. Borderline entries are retained for sensitivity checks but excluded from headline statistics.

**Coding rules for multi-vector campaigns.** We define *initial access (IA)* as the first successful foothold that enabled execution or credentialed access within the target environment. Subsequent steps in the same campaign (e.g., credential re-use, exploitation of known CVEs, lateral movement, or living-off-the-land) are recorded separately for context. Counts for step-types can therefore exceed the number of campaigns, and percentages are reported against the relevant resolved denominator (we state *N* in each caption; cf. the IA distribution figure).

**Disambiguation and tie-breaks.** When multiple sources describe near-simultaneous actions, IA is assigned by (i) earliest verifiable artefact time-stamp and, if tied, (ii) the analyst narrative that establishes precondition–consequence order. Ambiguous cases are marked IA-unclear and excluded from IA-percentage denominators for that break- down.

**Examples.** A phishing-led credential harvest followed by VPN token reuse is coded as IA=Social Engineering, with token reuse captured as a follow-on step. Phishing followed by exploitation of an already-known vulnerability (n-day) is IA=Social Engineering, while the n-day appears only in follow-on analysis rather than redefining IA.

### 3.3 Operationalizing the APT Definition

We use Definition 1 as both inclusion criterion and analytical frame. Each campaign is coded on four binary dimensions: **A**ctor type, **P**ersistence, **S**ophistication, **T**argeting, and included in the analytical subset if **A** = 1 and **P** + **S** + **T** ≥ 2. Observed evidence is mapped to Indicator of Compromise categories in [Hagen and Helkala, 2024] to preserve traceability from concept to data.

### 3.4 Data Collection and Validation

Sources are open and independently verifiable, including vendor and research threat intelligence reports, advisories, CSIRT publications, conference material, and verified news. Searches combine general engines with curated queries. AI systems are used only for lead discovery, never as direct inputs; all entries are validated against primary sources. After deduplication and alias resolution, **60** campaigns meet the APT definition and comprise the analytical subset.

### 3.5 Limitations

Public reporting can be incomplete or revised, potentially biasing persistence and targeting downward. Language and regional coverage can skew visibility. We mitigate through conservative thresholds, explicit attribution confidence recording, and release of the blinded archive for review and replication [Blinded, 2025].Results should be read as a documented snapshot of observable activity in the window, not an exhaustive census.

## 4. Results

This section answers RQ1–RQ4 using the analytical subset. Borderline records are retained for transparency in the blinded archive but excluded from quantitative statistics [Blinded, 2025]. Unless stated otherwise, percentages use the 60 denominator.

### 4.1 Geopolitical Attribution of APT Campaigns (RQ1)

National attribution follows the coding in §3.2 and the curated dataset [Blinded, 2025]. Table 1 reports the counts. Minor or contested attributions are merged into *Unknown/eCrime* per the study taxonomy.

**Table 1: Campaigns by attributed country or group and primary strategic focus. Minor attributions, for example India and South Korea, are merged into Unknown/eCrime**

| Country/Group | Number of Campaigns | Primary Target Focus (examples) |
|---|---|---|
| **Russia** | 17 | Ukraine and NATO/EU, diplomatic corps and webmail platforms, critical infrastructure and energy, destructive tooling and wipers |
| **China** | 16 | Government networks, telecom and network edge devices, Southeast Asia and Taiwan, EU governments, supply chain and identity abuse, botnets on end of life devices |
| **North Korea** | 15 | Cryptocurrency and financial sector, South Korean enterprises, software supply chains and developer ecosystems, defence and nuclear, technology companies |
| **Iran** | 9 | Israel and the broader Middle East, satellite and communications, oil & gas, aerospace and defence, spear phishing and impersonation campaigns |
| **Unknown/eCrime** | 3 | Enterprise intrusion and monetisation (for example FIN7, SCATTERED SPIDER, OldGremlin), cloud, SaaS, and retail; regional outliers |

### 4.2 Actor Sector Relationships (RQ2)

Actor sector variation concentrates in a few pairings. Figure 1 summarizes the eight core sectors used for cross actor comparisons. Additional sectors and long tail categories are available in the dataset [Blinded, 2025].

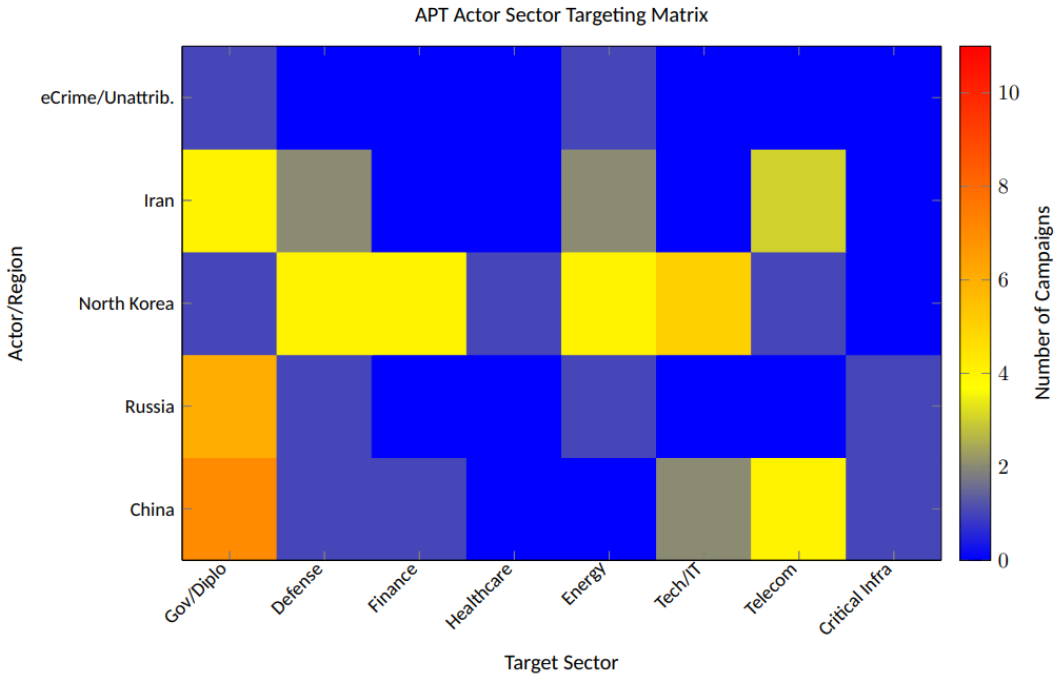


**Figure 1: Actor family by sector across confirmed campaigns ($N$ = 60; window=January 2024–July 2025). Cell counts reflect observed campaign–target pairs; totals can exceed $N$ where a campaign targeted multiple sectors or regions**

### 4.3 Initial Access Techniques (RQ3)

Initial access is led by a small set of vectors. Figure 2 shows campaign counts per vector.

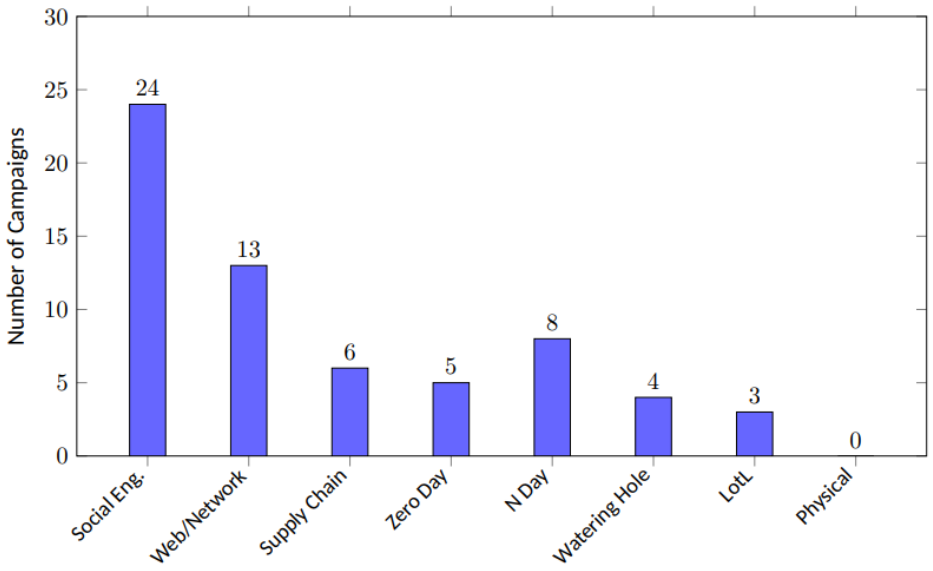


**Figure 2: Initial access distribution across confirmed campaigns ($N$ = 60; window=January 2024–July 2025). Social engineering dominates, while n-day reuse and zero-day exploitation appear less frequently; IA is defined as the first successful foothold, with follow-on steps recorded separately**

**Sequence-aware interpretation.** Actors frequently chain multiple steps after the first foothold (e.g., credential re-use, lateral movement, or exploitation of already-known vulnerabilities). We therefore treat *initial access* as the first successful foothold that triggers the earliest containment and triage decisions, while recording subsequent steps separately for context (cf. our IA distribution figure and campaign tables). This keeps the measure decision-relevant for defenders, without denying that later steps often escalate impact.

### 4.4 Detection Provenance (RQ4)

Most campaigns enter the public record via vendor threat reporting, followed by coordinated advisories and CSIRTs; tech vendors contribute a smaller share. Coattributions occur but do not change the majority pattern. organisation level tallies and copublication details are available in the curated archive [Blinded, 2025].

## 5. Discussion

**Table 2: Key metrics from the analytical subset**

| | |
|---|---|
| **Top actors (counts)** | Russia (17), China (16), North Korea (15), Iran (9), Unknown/eCrime (3). |
| **Top sectors (counts)** | Government/Diplomatic (19); Defense/Aerospace (8), Telecommunications (8), Energy/Nuclear (8). |
| **Initial access (counts)** | Social engineering (24), Web/Network (13), N day (8), Supply chain (6), Zero day (5). |

We interpret the analytical subset in light of RQ1–RQ4. .

### 5.1 RQ1: Geopolitical Actor Trends

Four primary state clusters, Russia (17), China (16), North Korea (15), and Iran (9), represent 95% of attributed campaigns, revealing shifts in operational tempo, targeting discipline, and collaborative behaviours. Russia's 17 campaigns reflect sustained operational tempo despite sanctions and exposed tooling. Persistence suggests acceptance of attribution costs or confidence that adaptation outpaces defensive responses. Distribution across APT28, APT29, and Sandworm maintains GRU and SVR separation with limited tactical convergence despite shared objectives. This compartmentalization provides resilience against comprehensive takedowns. China's 16 campaigns demonstrate industrial scale operations with uniform tradecraft, suggesting centralized coordination despite distributed execution. Consistent targeting of edge devices and identity systems indicates systematic intelligence requirements from central planning rather than opportunistic targeting. Absence of financially motivated operations reinforces purely strategic objectives. North Korea's 15 campaigns reveal operational maturation, with Lazarus sustaining concurrent campaigns across sectors while maintaining operational security. Documented

handoffs between access teams and ransomware operators suggest professionalized criminal partnerships blurring state and criminal boundaries. This hybrid model may emerge for sanctions constrained actors. Iran's 9 campaigns show increased sophistication in social engineering and targeting precision. Temporal clustering around regional events suggests specific collection requirements rather than persistent broad spectrum collection. Lower volume reflects capability constraints, selective targeting, or improved operational security. Minimal other nation state representation (3 Unknown/eCrime campaigns) suggests attribution challenges for emerging actors, reduced activity, or migration to gray zone operations below APT threshold. Concentration among established actors indicates high barriers to sustained APT operations, requiring technical capability, operational infrastructure, targeting intelligence, and sustained resourcing few actors can maintain.

### 5.2 RQ2: Sectoral Targeting Patterns

Government and diplomatic targets lead observed campaigns (19 instances), followed by defense, telecommunications, and energy sectors (8 each). The actor sector matrix exposes distinct strategic logics. China's concentration on telecommunications (4 campaigns) and government networks (7 campaigns) aligns with technology acquisition priorities and prepositioning for regional contingencies. Emphasis on network edge devices suggests systematic mapping of global routing infrastructure for persistent access enabling both peacetime collection and wartime disruption. North Korea's diverse targeting spans technology (5), defense (4), finance (4), and energy (4), reflecting dual pressures: sanctions driven revenue generation through cryptocurrency theft and strategic intelligence supporting nuclear and missile programs. Equal distribution across sectors suggests coordinated tasking across operational units with distinct missions, marking maturation from earlier predominantly financial opera- tions. Russia's government and diplomatic focus (6 campaigns) intensified following geopolitical tensions, shift- ing from broad spectrum targeting pre-2024 to concentrated NATO decision making infrastructure collection. Absence from finance and healthcare suggests improved operational security, strategic reorientation toward mil- itary intelligence priorities, or displacement of criminal activity to nonstate actors. Emphasis on destructive capa- bilities reinforces strategic rather than economic objectives. Iran's telecommunications emphasis (3 campaigns) correlates with regional surveillance objectives and bypassing international communications monitoring. Gov- ernment and diplomatic targeting (4 campaigns) shows temporal clustering around regional events, suggesting responsive rather than persistent collection. This reactive pattern differs from sustained Chinese and Russian approaches, potentially reflecting resource constraints or limited strategic objectives. These patterns indicate sector specific defensive investments must account for concentration and actor specific motivations. Identity and privilege management prove critical for telecommunications and energy facing Chinese and Russian inter- est. Transaction monitoring addresses DPRK financial targeting. Diplomatic communication hardening counters Russian and Iranian government focused operations. Recurrence of sector actor pairings across the 18 month window suggests persistent rather than episodic interest, requiring sustained defensive postures.

### 5.3 RQ3: Initial Access Emphasis

**Why initial access still matters.** Even when campaigns involve several chained techniques, the first foothold sets the time-critical actions: identity containment, scoping, and patch triage Mandiant [2025], CrowdStrike, Inc. [2025]. Later stages refine rather than replace these priorities, so reporting the initial vector remains a useful, decision-oriented signal for defenders, whereas later steps describe escalation paths.

Social engineering's dominance (24 campaigns, 40%) over technical exploitation reveals a shift in APT at- tack economics. Combined with web/network exploitation (13 campaigns) and N day usage (8 campaigns), these account for 75% of observed initial access, suggesting most APT objectives are achievable without zero day capabilities. Social engineering transcends simple phishing, encompassing sophisticated impersonation chains, compromise of trusted platforms, and exploitation of professional networking sites. Consistency across all major actors indicates convergent evolution toward exploiting human trust rather than technical vulnerabilities. This pattern is pronounced in diplomatic and government sectors, where spear phishing leverages geopolitical events to achieve compromise. Sustained use suggests human layer defenses lag behind technical controls. Web and network exploitation (13 campaigns) targets internet facing applications and VPN endpoints, emphasizing recently disclosed enterprise vulnerabilities. Exploitation within days of disclosure indicates APT actors maintain ready weaponization capability. This rapid operationalization challenges traditional patch windows and suggests the zero day versus N day distinction matters less than time to patch metrics. Low zero day usage (5 campaigns, 8.3%) contradicts assumptions about APT sophistication. Zero days appear reserved for high value targets where social engineering is infeasible or persistent access justifies operational cost. Selective deployment suggests rational resource allocation rather than capability constraints. Supply chain compromise (6 campaigns, 10%) enables

single intrusions to cascade across victims. Clustering among Chinese and Russian actors suggests this requires sustained infrastructure and victim research smaller actors cannot maintain. Targeting software pipelines and managed service providers indicates strategic multiplier selection. Limited watering hole (4 campaigns) and living off the land techniques (3 campaigns) as initial access may reflect detection bias, as these methods are harder to attribute. Their presence in sophisticated Chinese and Russian operations suggests these techniques require extensive reconnaissance and operational patience aligned with long term collection rather than immediate exploitation.

### 5.4 RQ4: Detection Provenance

Campaign discovery concentration through vendor threat reporting, followed by coordinated advisories and CSIRTs, reveals critical dependencies and gaps in the global detection ecosystem. Vendor dominance raises questions about visibility bias. Commercial security vendors' detection is limited to their deployment footprint, creating blind spots in regions with limited commercial tool adoption, sectors with legacy infrastructure, and organisations using open source or indigenous solutions. Geographical distribution of major vendors in North America, Western Europe, and developed Asian markets may result in underreporting of campaigns targeting Africa, Latin America, and Central Asia. This detection inequality could enable persistent adversary presence in undermonitored regions serving as staging grounds for broader campaigns. Coordinated advisories and CSIRTs' secondary role suggests na- tional and sector specific detection capabilities rarely achieve first detection of novel campaigns. This lag reflects resource constraints, limited cross organisational visibility, or defensive rather than hunting oriented missions. Few government first detections involved campaigns directly targeting government networks, suggesting victim proximity remains crucial for discovery. Minimal contribution from technology vendors indicates persistent chal- lenges in building security observability into products. Despite repeated targeting of edge devices, productivity suites, and development tools, these vendors rarely achieve first detection of abuse. Product telemetry remains optimized for operational rather than security purposes, missing early detection opportunities at platform level. Detection to disclosure timing varies by source. Vendor led discoveries typically reach public disclosure within 30-90 days, balancing responsible disclosure with commercial imperatives. Government led discoveries experi- ence longer delays, suggesting extended remediation or intelligence gain/loss calculations delaying attribution. These differences affect defensive response to emerging threats, with vendor driven disclosure providing more actionable intelligence in tactical timeframes. Absence of victim self reporting suggests organisations lack detec- tion capabilities or reputational and regulatory concerns discourage disclosure. This silence limits understanding of dwell time, impact, and defensive effectiveness, creating an incomplete threat picture that may overemphasize detected campaigns while missing persistent, undetected intrusions.

**Provenance caveats.** First-sighting claims are not uniformly distributed across sensors or regions; they depend on product footprint, customer mix, and disclosure practices. As a result, "who saw it first" is partly a function of market presence rather than global prevalence. To reduce this bias, we triangulate across multiple report types (vendor advisories, CERT bulletins, and community write-ups) and time-stamp the earliest public disclosure per campaign, while keeping later confirmations for context.

### 5.5 Synthesis: Cross-Cutting Implications for Attribution and Defence

**Where conventional wisdom comes from, and where it fails.** Much of the field's narrative is produced by vendors and agencies with uneven telemetry and staggered disclosure timelines. This can over-amplify techniques that their sensors capture well and understate activity outside their footprint. Baseline guidance that treats threats uniformly is useful, but our evidence indicates better outcomes when priorities are actor- and sector- specific: identity hardening where social engineering dominates, accelerated patch triage where nday reuse is common, and targeted monitoring where livingofftheland is typical National Cyber Security Centre [2024], Cy- bersecurity and Infrastructure Security Agency [2025].

Integrated analysis of actor patterns, targeting preferences, access methods, and detection dynamics reveals systemic implications requiring reconsideration of attribution methodologies and defensive architectures.

Attribution confidence emerges as hierarchical. National clustering provides strongest signal, reinforced by predictable targeting patterns and tooling preferences within geopolitical groups. Group specific tooling chains like MISTPEN for Lazarus or custom loaders for APT29 provide secondary confirmation when combined with infrastructure overlaps. However, widespread commodity framework adoption across actors, Cobalt Strike appearing in all major nation state campaigns, diminishes attribution value. This hierarchy suggests attribution requires differential weighting: infrastructure patterns and bespoke malware carry substantial weight, commodity tools

provide minimal value, and targeting patterns serve as corroborating rather than primary evidence. organisations must rapidly assess whether activity represents targeted state interest or opportunistic criminal activity.

Defensive architectures must evolve beyond uniform controls to account for pronounced actor sector convergence. Chinese telecommunications concentration, Russian government focus, and North Korean financial emphasis enable risk calibrated implementation. While identity management emerges as universal given social engineering dominance, implementation must reflect threat specific adaptations. Telecommunications facing Chinese interest require supply chain verification and hardware attestation for edge devices. Financial institutions under North Korean targeting need transaction anomaly detection emphasizing cryptocurrency conversion and cross border transfers. Government entities confronting Russians must implement diplomatic communication isolation with air gapped authentication resistant to destructive attacks. Sector specific adaptations grounded in observed preferences yield superior outcomes compared to generic frameworks.

Vendor centric detection creates systematic gaps sophisticated actors exploit for persistent access. organisations in regions with limited commercial deployment, sectors using legacy infrastructure, or environments where living off the land predominates must assume APTs operate below detection thresholds. This necessitates shifting from detection centric models assuming timely discovery to resilience architectures assuming persistent presence. Such architectures prioritise segmentation limiting lateral movement, immutable backups for restoration, and deception technologies increasing adversary costs. Hunt operations should focus on limited vendor visibility areas: operational technology networks, non Windows environments, and legacy systems where evidence suggests persistent footholds.

Temporal dynamics require defensive strategies operating across multiple horizons simultaneously. Rapid vulnerability weaponization with N day exploitation within days demands immediate patch protocols compressing testing windows. Sustained social engineering effectiveness indicates human defenses require continuous reinforcement. The 18 month persistence of actor sector relationships demonstrates sustained strategic interest requiring continuous evolution. This complexity requires programs maintaining simultaneous focus on immediate vulnerability response, sustained awareness training, and long term transformation, with resource allocation reflecting relative risk.

These findings challenge conventional wisdom treating threats uniformly and assuming technical controls adequately defend against nation states. Evidence supports differentiated models where attribution drives response, sector determines prioritization, visibility gaps guide hunting, and temporal dynamics shape planning. organisations facing APT threats must develop adaptive programs reflecting sector, geography, and specific adversaries likely targeting their assets, with investments calibrated to observed rather than theoretical behaviours.

### 5.6 Network Dynamics and Collaboration Patterns

Network analysis reveals distinct operational patterns across actor groups. North Korean operations demonstrate cohesion with limited cross actor linkages. Two documented bridges underpin the DPRK subgraph: Andariel's initial access handoff to Play ransomware, and the Lazarus UNC2970 linkage via MISTPEN chain. Beyond these, overlaps remain at tradecraft level rather than direct infrastructure or tooling sharing. Russian operations exhibit compartmentalization across intelligence services. APT28, APT29, and Sandworm show limited evidence of direct implant or command and control infrastructure sharing, with similarities confined to thematic lures and targeting patterns. This separation aligns with documented interservice boundaries between GRU and SVR, suggesting organisational silos persist even when objectives converge. Chinese activity manifests through parallel operational cells emphasizing telecommunications infrastructure and identity systems. Clusters concentrate on network edge devices and supply chain compromise, developing bespoke malware within distinct cells while demonstrating limited intercluster tool reuse in public campaigns. This pattern suggests coordinated strategic direction with distributed tactical execution. Widespread commodity framework presence across actors underscores an attribution principle: shared tooling does not constitute collaboration evidence. Frameworks like Cobalt Strike appear across multiple nation state operations and must be excluded as adjacency evidence. Only group specific tooling chains with documented operational linkages, such as MISTPEN between Lazarus and UNC2970, provide valid actor relationship evidence.

**Decision-oriented implications.** Given the IA distribution in Fig. 2, prioritise identity-centric containment over rare zero-day scenarios. Because re-use of known vulnerabilities materially exceeds zero-day use in our window, patch triage against the currently exploited set should outrank generic "patch everything" drives. Use the actor×sector heatmap (Fig. 1) to add a thin, sector-specific overlay to the baseline rather than treating all threats

uniformly. Finally, since first-sighting is partly a function of telemetry footprint, require multi-source corroboration before rebalancing investments on the basis of a single “who saw it first” claim.

## 6. Conclusion

We analyzed 60 verified APT campaigns from January 2024 to July 2025, revealing a paradoxical threat landscape. APTs deploy zero day exploits while relying on decades old social engineering, maintain complex infrastructure while exploiting basic configuration errors, and pursue strategic nation state objectives while adapting opportunistically. This duality challenges conventional models positioning APTs as uniformly sophisticated, instead revealing actors who select minimum viable techniques to achieve objectives. Activity concentration among four nation state clusters, Russia (17), China (16), North Korea (15), and Iran (9), demonstrates tight coupling between geopolitical tensions and cyber operations. Within this state directed framework, we observe pragmatic adaptations driven by resource constraints, operational security, and target availability. North Korea’s dual focus on revenue and strategic intelligence exemplifies tension between state direction and operational autonomy. Documented handoffs between nation state operators and criminal ransomware groups blur traditional boundaries, suggesting APTs operate across a spectrum rather than within rigid frameworks. For defenders, this complexity strains organisational capabilities. Campaigns exploited vulnerabilities from firmware implants to social manipulation, demanding expertise across widening technology arrays. organisations in telecommunications and government face compound challenges where measures optimized for one adversary prove inadequate against another. Rapid vulnerability weaponization alongside persistent legacy techniques requires programs operating across multiple horizons, maintaining immediate response and long term transformation simultaneously. Systematic visibility gaps from detection provenance analysis reveal fragmented, incomplete defensive capabilities. Detection concentration in commercial vendors creates blind spots sophisticated actors exploit, while limited technology vendor contribution suggests persistent disconnects between product development and security requirements. Gaps extend beyond technical capabilities to expertise, where the skill gap widens despite increased investment. Evidence suggests individual organisational responses cannot adequately address APT threats leveraging global infrastructure and cross border flexibility. These findings compel reconsideration from isolated organisational responses to collaborative ecosystem approaches. Cross sector sharing must evolve beyond indicators to include behavioural patterns, attribution methodologies, and defensive innovations raising collective resilience. International cooperation becomes essential for real time defensive coordination matching transnational adversary tempo. Actor sector persistence indicates targeted sectors share defensive challenges benefiting from collaboratively developed rather than duplicated frameworks. The paradoxical APT nature, advanced yet basic, strategic yet opportunistic, persistent yet adaptive, reflects fundamental asymmetry where attackers need single successes while defenders maintain continuous vigilance across expanding surfaces. This asymmetry requires systemic approaches acknowledging geopolitical drivers and operational realities of defending interconnected systems. As APTs evolve tradecraft responding to improvements and attribution, collaborative, adaptive, strategically informed approaches become essential for organisational and national security.

## 7. Future Work

Extending the observation window to multiple years would enable longitudinal analysis of campaign evolution, ac- tor adaptation, and correlation with geopolitical events. A five year dataset could capture cyclical patterns tied to political transitions or economic cycles invisible in the current 18 month snapshot. Geographic parameters including attack infrastructure geolocation, operational hours, and network routing paths would enable spatial temporal pattern recognition revealing operational security practices. Technical telemetry at packet and behavioural levels would support machine learning approaches for attribution. Integration of sanctions regimes, diplomatic ten- sions, and military readiness indicators could reveal predictive relationships between real world events and cyber operations. The actor relationship edge list provides foundation for network analysis the current dataset cannot support. With expanded coverage, graph neural networks trained on actor infrastructure target relationships could reveal latent patterns. Time series analysis might identify critical transitions preceding campaign intensifi- cation or collaboration. Predictive modeling could transform defense from reactive to anticipatory. Classification models could assess attribution probability given partial observables. Survival analysis could model time to de- tection based on sector, actor, and technique combinations. Agent based simulations with observed behaviours could explore intervention strategies. Incorporating non English sources and regional CERT partnerships would reduce Western bias and surface invisible campaigns. Standardized coding schemas with inter rater reliability protocols would enable distributed collection while maintaining quality. The objective extends beyond analysis to operational impact. A continuously updated dataset with rich attributes would enable evidence based strate- gies, quantitative risk

assessment, and optimal resource allocation. Transition from static analysis to dynamic prediction would fundamentally alter how organisations respond to nation state threats, moving from anecdotal intelligence to data driven strategies grounded in empirical adversary behaviour.

**Declaration of generative AI and AI-assisted technologies in the writing process**: The author used AI assisted tools to support language editing and formatting: (i) improving clarity, grammar, and idiom; (ii) reducing dyslexia related spelling/ordering errors; and (iii) troubleshooting LATEX minutiae (bibliogra- phy keys, table widths, listing frames). All AI suggested edits were reviewed and approved by the author. Any remaining typos are proudly human.

**Ethics declaration**: This research uses exclusively publicly available open source data, requiring no ethical clearance. The dataset is anonymized for review to protect author identities, consistent with double blind review protocols.

## References


Blake Strom, Andy Applebaum, Douglas Miller, Kathryn Nickels, Adam Pennington, and Cody Thomas. Mitre att&ck®: Design and philosophy. Technical Report PRS 19-01075-28, The MITRE Corporation, March 2020. URL https://attack.mitre.org/docs/ATTACK_Design_and_Philosophy_March_2020.pdf.

Brandon Valeriano, Benjamin Jensen, and Ryan C. Maness. *Cyber Strategy: The Evolving Character of Power and Coercion*. Oxford University Press, 2018. ISBN 978-0-19-061809-4.

CrowdStrike, Inc. 2025 global threat report, February 2025. URL https://www.crowdstrike.com/en-us/ global-threat-report/.

Cybersecurity and Infrastructure Security Agency. Cross-sector cybersecurity performance goals (cpgs), Octo- ber 2023. URL https://www.cisa.gov/cross-sector-cybersecurity-performance-goals. Baseline practices applicable across critical infrastructure.

Cybersecurity and Infrastructure Security Agency. Known exploited vulnerabilities (kev) catalog, 2025. URL https://www.cisa.gov/known-exploited-vulnerabilities-catalog. Live catalogue, updated con- tinuously.

Dataset Blinded. Apt dataset 2024–2025, 2025. URL https://osf.io/bhv53/?view_only=.

Eric M. Hutchins, Michael J. Cloppert, and Rohan M. Amin. Intelligence-driven computer network defense in- formed by analysis of adversary campaigns and intrusion kill chains. In *Proceedings of the 6th International Conference on Information Warfare and Security (ICIW)*, pages 113–125. Academic Conferences International Limited, 2011.

European Union Agency for Cybersecurity (ENISA). Enisa threat landscape 2025, October 2025. URL https://www.enisa.europa.eu/sites/default/files/2025-10/ENISA%20Threat%20Landscape%202025.pdf. TLP:CLEAR.

Information technology — security techniques — information security incident management — part 1: Principles of incident management, 2023. URL https://www.iso.org/standard/78973.html.

Mandiant. M-trends 2025 report, April 2025. URL https://services.google.com/fh/files/misc/m-trends-2025-en.pdf.

Marcus Willett. Lessons of the solarwinds hack. *Survival*, 63(2):7–26, 2021. doi: 10.1080/00396338.2021. 1906001. URL https://www.tandfonline.com/doi/full/10.1080/00396338.2021.1906001.

National Cyber Security Centre. Recommended types of mfa, September 2024. URL https://www.ncsc.gov.uk/collection/mfa-for-your-corporate-online-services/recommended-types-of-mfa. Guidance for corporate online services.

National Institute of Standards and Technology. The nist cybersecurity framework (csf) 2.0. NIST Cybersecurity White Paper NIST CSWP 29, NIST, February 2024. URL https://doi.org/10.6028/NIST.CSWP.29.

Peter Chen, Lieven Desmet, and Christophe Huygens. A study of advanced persistent threats. *IFIP In- ternational Conference on Communications and Multimedia Security*, pages 63–72, 2014. doi: 10.1007/ 978-3-662-44885-4 5.

Ralph Langner. Stuxnet: Dissecting a cyberwarfare weapon. *IEEE Security & Privacy*, 9(3):49–51, 2011. doi: 10.1109/MSP.2011.67.

Raymond Andre´ Hagen and Kirsi Helkala. The complexity of contemporary indicators of compromise. In *Pro- ceedings of the 23rd European Conference on Cyber Warfare and Security (ECCWS 2024)*, pages 693–703, Reading, UK, June 2024. Academic Conferences International Limited. doi: 10.34190/eccws.23.1.2149. URL https://papers.academic-conferences.org/index.php/eccws/article/view/2149.

Thomas Rid. *Cyber War Will Not Take Place*. Oxford University Press, 2013. ISBN 978-0-19-933063-8.