Advanced International Journal for Research (AIJFR)
E-ISSN: 3048-7641 ● Website: www.aijfr.com ● Email: editor@aijfr.com

The Evolution of Attack Techniques and Their Growing Complexity: A Longitudinal Analysis of MITRE ATT&CK, 2020–2026, and Its Implications for Detection

Jini Fernandez

Independent Researcher, Riyadh, Saudi Arabia

DOI: https://doi.org/10.63363/aijfr.2026.v07i06.8884

Abstract
Attack techniques have changed from self-propagating worms to targeted, multi-stage campaigns that abuse legitimate tools, identities, cloud services and software supply chains. Most accounts of this change are narrative. This paper combines a narrative review with a longitudinal measurement of MITRE ATT&CK, the most widely used public knowledge base of adversary behavior. Using the official release notes of thirteen Enterprise ATT&CK releases from July 2020 (v7) to April 2026 (v19), we measure how documented adversary behavior grew. The number of documented techniques and sub-techniques rose from 428 to 697 (+63%), with sub-techniques growing faster than techniques (+75% versus +42%). From October 2023 to April 2026, documented campaigns more than doubled, from 23 to 56. The framework also expanded to new platforms, including network devices, containers and software-as-a-service applications. We then trace six shifts in attack vectors: mass-propagating malware, targeted advanced persistent threats, ransomware as a criminal business, living off the land, identity and supply-chain compromise, and AI-assisted operations. We argue that complexity has grown less through new attack categories than through finer variation within known behaviors, more platforms, and greater use of legitimate tools. These trends shift detection away from signatures toward behavioral, context-rich and continuously maintained detection engineering, and we discuss implications for SIEM and security operations practice.

Keywords: Attack Techniques, MITRE ATT&CK, Threat Landscape, Advanced Persistent Threats, Ransomware, Living off the Land, Supply Chain Attacks, Detection Engineering

1. Introduction

Defenders often say attacks have become more complex, but the claim is rarely measured. This paper measures one observable form of that complexity, the growth of documented adversary behavior, and connects it to how attack vectors have shifted over nearly four decades.

MITRE ATT&CK has become the shared vocabulary for describing what adversaries do after, and increasingly before, they gain access to a system [1]. It organizes behavior into tactics (the adversary's

goals), techniques (how a goal is achieved) and sub-techniques (more specific variants), and links them to documented threat groups, software and campaigns. A recent survey describes its wide adoption across threat intelligence, detection, emulation and assessment [2]. Because ATT&CK is versioned and its release notes report counts for each release, its growth provides a consistent, public record of how much adversary behavior defenders must account for.

Existing surveys describe the evolution of particular threat classes, such as advanced persistent threats [3], ransomware [4] and software supply-chain attacks [5]. Fewer works look across these classes, and fewer still quantify the change. At the same time, research shows that detection struggles to keep pace: widely used detection rules can be evaded [6], and claimed coverage of ATT&CK techniques does not consistently mean the same detection capability [7].

This paper asks three questions:

- **RQ1.** How has documented adversary behavior in ATT&CK grown across releases from 2020 to 2026?
- **RQ2.** Which shifts in attack vectors explain that growth?
- **RQ3.** What do these shifts imply for detection engineering and security operations?

The contributions are: (1) a longitudinal dataset of Enterprise ATT&CK counts across thirteen releases, each traced to its official release notes; (2) an analysis of where growth occurred, separating structural changes in the framework from growth in documented behavior; (3) a synthesis of six major shifts in attack vectors; and (4) implications for detection practice.

Section 2 reviews background and related work. Section 3 describes the method. Section 4 presents the ATT&CK measurements, and Section 5 traces the shifts in attack vectors. Section 6 discusses implications for detection, and Section 7 covers limitations and conclusions.

2. Background and Related Work

2.1. Models of Adversary Behavior

The intrusion kill chain described an intrusion as a sequence of stages, from reconnaissance and weaponization to actions on objectives, and argued that defenders can disrupt an attack at any stage [8]. SIEM research has since applied kill-chain stages to the analysis of security logs [9]. ATT&CK extended the idea with a much finer description of post-compromise behavior, grounded in observed intrusions rather than theoretical attack paths [1]. It grew out of MITRE's research on detecting post-compromise adversary behavior in enterprise networks [1]. Since then it has been extended to new platforms, given sub-techniques, and linked to threat groups, software and named campaigns [2].

2.2. Measurements of the Threat Landscape

Other public sources measure different aspects of attacks. Incident-response data shows which initial access vectors succeed: Mandiant reported that exploits were the most common initial infection vector in its 2024 investigations at 33%, followed by stolen credentials at 16%, and that the global median dwell time was 11 days [10]. Government catalogs track which vulnerabilities attackers actually use; the US

Known Exploited Vulnerabilities catalog was created in November 2021 to drive remediation of vulnerabilities under active exploitation [11]. Threat-intelligence sharing exchanges indicators and reports between organizations, although with uneven quality and timeliness [12]. These sources measure frequency and impact. ATT&CK, by contrast, measures breadth: how many distinct behaviors defenders need to recognize.

2.3. Surveys of Attack Evolution

Surveys have traced the evolution of specific threats. Alshamrani et al. reviewed the techniques and life cycle of advanced persistent threats [3]. Oz et al. traced ransomware from early families to ransomware-as-a-service and multi-extortion [4]. Ladisa et al. built a taxonomy of attacks on open-source software supply chains [5]. These works are thorough within their scope, but they do not compare threat classes against a common, quantified baseline. Studies of ATT&CK use have examined how vendors map detections to techniques [7] and how the framework is applied in research and industry [2], but not how the framework itself has grown.

3. Method

The study combines a longitudinal measurement of ATT&CK with a narrative review of attack-vector shifts.

3.1. ATT&CK Data

For each Enterprise ATT&CK release from v7 (July 2020) to v19 (April 2026), we recorded the counts of tactics, techniques and sub-techniques reported in the official release notes [13–25]. For releases from v14 (October 2023) onward, we also recorded Enterprise counts of groups, software and campaigns, which the release notes report consistently for this period. Earlier releases report some of these counts only for all domains combined or only as additions, so they were excluded from those series to avoid mixing scopes.

The series starts at v7 because that release introduced sub-techniques and restructured the framework: many former techniques became sub-techniques, and the number of Enterprise techniques fell from 266 to 156 [13]. Counts before v7 are therefore not comparable with later ones.

3.2. Metrics

The main metric is **documented behaviors**, defined as techniques plus sub-techniques in a release. It measures the number of distinct adversary behaviors that a defender would need to recognize to have complete conceptual coverage. Secondary metrics are the share of behaviors that are sub-techniques, which indicates refinement within known behaviors, and the counts of groups, software and campaigns, which indicate how much observed activity the framework documents.

Growth in ATT&CK reflects both adversary change and MITRE's own decisions about scope and structure. We therefore identify structural changes separately, such as new tactics, merged domains and new platforms, using the release notes.

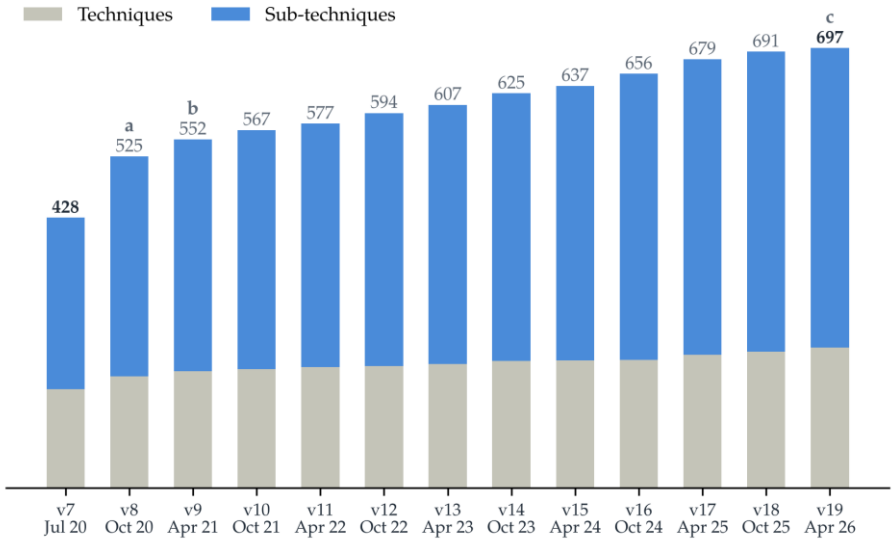
3.3. Review of Attack-Vector Shifts

The narrative review covers foundational analyses of major incidents, peer-reviewed surveys, and government and industry reports. Sources were selected for their direct relevance to a change in how attackers gain access, move, persist or achieve their objectives.

4. Results: The Growth of Documented Adversary Behavior

Documented adversary behavior in Enterprise ATT&CK grew by 63% between July 2020 and April 2026, from 428 to 697 techniques and sub-techniques (Figure 1) [13–25].

Figure 1: Growth of Enterprise ATT&CK Techniques and Sub-Techniques, v7 (July 2020) to v19 (April 2026)



a PRE-ATT&CK merged into Enterprise; Reconnaissance and Resource Development tactics and a Network platform added
b Containers and Google Workspace platforms added; cloud provider platforms combined into IaaS
c Defense Evasion split into Stealth and Defense Impairment, giving 15 tactics

Source: official Enterprise ATT&CK release notes, v7–v19 [13–25].

4.1. Steady Growth, Mostly Within Known Behaviors

Growth was steady rather than episodic. The largest single increase, 97 behaviors at v8, was structural: PRE-ATT&CK was merged into Enterprise, adding the Reconnaissance and Resource Development tactics [14]. After v8, each release added between 6 and 27 behaviors.

Most new behavior was added as variants of existing techniques rather than as new techniques. Of the 269 behaviors added between v7 and v19, 203 were sub-techniques and 66 were techniques. Sub-techniques grew by 75% and techniques by 42%, and the sub-technique share of all behaviors rose from 64% to 68%. This pattern suggests that adversaries more often find new ways to achieve familiar goals than invent new categories of attack. For defenders, it means that a detection written for one variant of a technique may miss its siblings.

4.2. More Documented Activity

For the period in which Enterprise counts are reported consistently, documented activity grew faster than documented behavior (Table 1) [20–25].

Table 1: Enterprise ATT&CK Groups, Software and Campaigns, v14–v19

Release	Date	Groups	Software	Campaigns
v14	Oct 2023	141	648	23
v15	Apr 2024	148	677	28
v16	Oct 2024	159	710	34
v17	Apr 2025	166	755	47
v18	Oct 2025	172	784	52
v19	Apr 2026	174	821	56

Campaigns grew by 143%, software by 27% and groups by 23%. Campaigns were introduced as an object type only in v12 [18], so their rapid growth partly reflects MITRE documenting earlier activity in the new format, not only new attacks.

4.3. More Platforms and a Stronger Defensive Layer

ATT&CK's scope widened to follow attackers into new environments. v8 added a Network platform for switches, routers and other network devices [14], and v9 added Containers and Google Workspace and combined the separate cloud-provider platforms into a single IaaS platform [15]. In v18, MITRE replaced the detection guidance inside techniques with structured Detection Strategies and Analytics, listing 691 Enterprise detection strategies, matching the 691 techniques and sub-techniques in that release, and 1,739 analytics [24]. In v19, the Defense Evasion tactic was split into Stealth and Defense Impairment, raising the number of Enterprise tactics to 15 [25]. The framework has therefore grown in two directions: more attacker behavior to recognize, and more structured guidance on how to detect it.

5. How Attack Vectors Shifted

The growth measured in Section 4 is the latest stage of a longer change: attacks moved from noisy, indiscriminate malware toward quiet, targeted operations that blend into legitimate activity. Table 2 summarizes six overlapping shifts.

Table 2: Six Shifts in Attack Vectors

Shift	Approximate period	Typical way in	What changed for defenders
Mass-propagating malware	Late 1980s–2000s	Self-spreading exploits of network services	Speed: infections spread faster than people could respond
Targeted persistent intrusions	2010s onward	Spear-phishing, custom malware, zero-days	Stealth and persistence over months
Ransomware as a business	Mid-2010s onward	Bought access, phishing, exposed services	Specialized criminal roles and multi-extortion
Living off the land	2010s onward	Built-in administrative tools after initial access	Malicious activity looks like administration
Identity and supply-chain compromise	Late 2010s onward	Stolen credentials, trusted software and providers	The trusted path itself is the attack path
AI-assisted operations	Mid-2020s onward	Automated exploitation and reconnaissance	Scale and speed of skilled work increase

5.1. Mass-Propagating Malware

The 1988 Morris worm exploited weaknesses in common network services to spread across the early Internet [26]. Later worms showed how fast this could happen: Code-Red spread across vulnerable web servers in 2001 [27], and in 2003 the Slammer worm infected more than 90% of vulnerable hosts within 10 minutes, at least 75,000 hosts in total [28]. These attacks were loud and indiscriminate. Intrusion detection research had long distinguished between detecting deviations from normal behavior [29] and matching known attack patterns [30]; against fast worms, patching, perimeter filtering and signatures for known malicious code were the practical defense.

5.2. Targeted Persistent Intrusions

Attackers then moved from spreading widely to reaching specific targets. Advanced persistent threats combine reconnaissance, tailored initial access, persistence and slow lateral movement over long periods [3]. Such campaigns could rely on vulnerabilities unknown to defenders: a study of real-world data identified 18 vulnerabilities exploited before public disclosure and found that a typical zero-day attack lasted 312 days on average [31]. Stuxnet demonstrated that a targeted operation could reach isolated industrial systems and cause physical effects [32]. The intrusion kill chain was developed in this period to help defenders reason about such multi-stage campaigns [8]. The defender's problem shifted from blocking known malware to finding a patient, adaptive adversary.

5.3. Ransomware as a Criminal Business

Ransomware evolved from isolated malware into an organized economy. An analysis of 1,359 samples from 15 families observed between 2006 and 2014 found that most relied on relatively simple locking, encryption or deletion techniques [33]. Ransomware-as-a-service, which emerged in 2015, then let developers rent their tools to affiliates who carry out intrusions [4]. Some families added double extortion, demanding payment both to decrypt data and to prevent stolen data from being published [4]. The result was specialization, with different actors performing different parts of an attack, which makes intrusions more professional and repeatable.

5.4. Living off the Land

Attackers increasingly use tools already present on the victim's systems. A systematic study of Windows malware analyzed how such living-off-the-land techniques are used in the wild [34]. In 2024, CISA, NSA, FBI and partner agencies from Australia, Canada, the United Kingdom and New Zealand issued joint guidance warning that these techniques are hard to detect: they abuse trusted administrative tools, leave few conventional indicators, and hide in large volumes of legitimate activity [35].

5.5. Identity and Supply-Chain Compromise

As organizations moved to cloud services and single sign-on, identities and trusted suppliers became attack paths. In December 2020, the US government ordered federal agencies to disconnect affected SolarWinds Orion products after attackers compromised the product's code [36]. Open-source ecosystems became another route: a review of 174 malicious packages distributed through npm, PyPI and RubyGems between 2015 and 2019 showed how attackers inject code into dependency trees [37], and attacks on open-source supply chains now follow many distinct paths, from malicious packages to compromised build systems [5]. Incident-response data shows the same shift at initial access: stolen credentials were the second most common initial infection vector in Mandiant's 2024 investigations,

behind exploits [10]. ATT&CK's own campaign records include supply-chain operations such as the 3CX compromise, added in v18 [24].

5.6. AI-Assisted Operations

Large language models have begun to lower the skill and time needed for parts of an attack. Surveys document a fast-growing body of work on large language models in security, on both offense and defense [38, 39]. In a 2024 study, a GPT-4-based agent exploited 87% of a set of 15 one-day vulnerabilities when given their CVE descriptions, but only 7% without them [40]. In November 2025, Anthropic reported disrupting a campaign in which a state-sponsored group used an AI coding agent to perform an estimated 80–90% of the work in attempted intrusions against about thirty organizations [41]. AI also creates new attack surface: code assistants can produce insecure code [42], and code-generating models frequently recommend software packages that do not exist, names that attackers could register to distribute malicious code [43]. Evidence of AI-enabled attacks is still limited and largely comes from vendors and researchers, but it points to faster and more scalable operations.

6. Implications for Detection

The shifts above move detection away from recognizing known malicious artifacts toward recognizing suspicious behavior in context. Five implications follow for SIEM and security operations practice, whose role is to turn this telemetry into actionable alerts [44].

From indicators to behavior. Signatures and indicators worked against mass malware, but living-off-the-land activity changes file names, arguments and hashes and leaves few fixed indicators [35], and technical indicators in threat-intelligence feeds have short lifetimes because adversaries change infrastructure quickly [45]. Detections must describe behavior, such as an unusual parent process, an administrative tool run by an unexpected account, or a rare sequence of actions. Anomaly-detection methods can support this [46], although they behave differently in security than in other domains because errors are costly and representative attack data is scarce [47]. Platform-neutral behavioral rule formats such as Sigma help share detection logic [48], but behavioral rules can still be evaded when written too narrowly [6].

Coverage is a moving target. With 697 documented behaviors in Enterprise ATT&CK and growth in every release since v7 (Section 4), coverage measured against a fixed list is out of date within months. Sub-technique growth makes this sharper: a rule mapped to a technique may cover only one of its variants. Studies of endpoint products show that claimed coverage of a technique does not consistently mean the same detection capability [7], and studies of production rulesets show that a small fraction of rules generates most alerts [49] and that rule design choices strongly affect noise [50]. Coverage should be reviewed at each ATT&CK release and measured by tested detections, not by mappings.

Correlation across stages. Multi-stage intrusions are rarely visible in a single event. Alert-correlation research has long sought to combine low-level alerts into descriptions of attacks [51] and to reconstruct attack scenarios from related alerts [52]. More recent systems correlate suspicious information flows into attack stages [53] or use provenance to rank alerts [54]. This matters because attacks are rare compared with normal activity, so detectors with even small false-alarm rates overwhelm analysts [55]. In one SOC, only about 0.01% of network alerts over four years were associated with true attacks [56]; analysts

report that false positives dominate their queues [57], and SOC studies link this to operational friction [58] and burnout [59].

Identity and cloud telemetry. When credentials and trusted suppliers are the way in, the most important evidence sits in identity-provider logs, cloud control planes and software-delivery systems rather than on endpoints alone. Research has shown the value of mining heterogeneous enterprise logs for suspicious behavior [60] and of modeling internal login paths to detect lateral movement [61]. A global median dwell time of 11 days [10] leaves little room for telemetry gaps in these sources.

Detection engineering as a continuous process. ATT&CK now ships structured detection strategies and analytics alongside techniques [24], and its twice-yearly releases add new behaviors each time. Detection content therefore needs the same life cycle as software: version control, testing against emulated techniques, scheduled review, and retirement. AI-assisted attacks, if they grow as early evidence suggests [40, 41], will shorten the time defenders have to update detections after a new technique appears. Defenders are adopting AI as well: SOC analysts already use LLMs as on-demand aids [62], but LLM-based alert triage has been found to miss a substantial share of attack-related alerts [63], so AI assistance does not remove the need for engineered detections.

7. Limitations and Conclusions

7.1. Limitations

- **ATT&CK measures documentation, not attacks.** Its growth reflects what MITRE and its contributors observed, reported and chose to add, as well as structural decisions such as merging domains or splitting tactics. Counts therefore indicate the breadth of known behavior, not the frequency or impact of attacks.
- **Reporting bias.** Public threat reporting over-represents some regions, sectors and adversaries. Behaviors seen mostly in unreported incidents are under-counted.
- **Limited time span.** Comparable counts begin only at v7 (July 2020), and consistent Enterprise counts of groups, software and campaigns only at v14 (October 2023).
- **Narrative evidence.** The review of attack-vector shifts in Section 5 is qualitative and draws partly on government and industry reports.

7.2. Conclusions

Attacks have become more complex in a measurable way. Between July 2020 and April 2026, documented adversary behavior in Enterprise ATT&CK grew by 63%, mostly through new variants of existing techniques, and documented campaigns more than doubled in two and a half years. Over a longer span, attack vectors shifted from mass-propagating malware toward targeted operations that abuse legitimate tools, identities and trusted suppliers, with early signs that AI is increasing their speed and scale. For defenders, the consequence is that detection can no longer rely on recognizing known malicious artifacts. It must describe behavior in context, use identity and cloud telemetry, and be maintained continuously as the documented threat landscape grows. Future work could extend the measurement to the Mobile and ICS domains and link ATT&CK growth to incident frequency data.

Data Availability

The data analyzed in this study are publicly available in the MITRE ATT&CK release notes cited as references [13–25].

Conflict of Interest

The author declares no conflicts of interest.

Acknowledgement

This research received no external funding. During the preparation of this manuscript, the author used Claude (Anthropic) for the purposes of drafting text, structuring the paper, and identifying and verifying references, collecting counts from public release notes. The author has reviewed and edited the output and takes full responsibility for the content of this publication.

References

1. B.E. Strom, A. Applebaum, D.P. Miller, K.C. Nickels, A.G. Pennington, C.B. Thomas, “MITRE ATT&CK: Design and Philosophy”, Technical Report, The MITRE Corporation: McLean, VA, USA, 2018.
2. B. Al-Sada, A. Sadighian, G. Oligeri, “MITRE ATT&CK: State of the art and way forward”, ACM Computing Surveys, 2025, 57, 12. https://doi.org/10.1145/3687300
3. A. Alshamrani, S. Myneni, A. Chowdhary, D. Huang, “A survey on advanced persistent threats: Techniques, solutions, challenges, and research opportunities”, IEEE Communications Surveys & Tutorials, 2019, 21, 1851–1877. https://doi.org/10.1109/COMST.2019.2891891
4. H. Oz, A. Aris, A. Levi, A.S. Uluagac, “A survey on ransomware: Evolution, taxonomy, and defense solutions”, ACM Computing Surveys, 2022, 54, 238. https://doi.org/10.1145/3514229
5. P. Ladisa, H. Plate, M. Martinez, O. Barais, “SoK: Taxonomy of attacks on open-source software supply chains”, Proceedings of the 2023 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 2023, pp. 1509–1526. https://doi.org/10.1109/SP46215.2023.10179304
6. R. Uetz, M. Herzog, L. Hackländer, S. Schwarz, M. Henze, “You cannot escape me: Detecting evasions of SIEM rules in enterprise networks”, Proceedings of the 33rd USENIX Security Symposium, Philadelphia, PA, USA, 2024, pp. 5179–5196.
7. A. Virkud, M.A. Inam, A. Riddle, J. Liu, G. Wang, A. Bates, “How does endpoint detection use the MITRE ATT&CK framework?”, Proceedings of the 33rd USENIX Security Symposium, Philadelphia, PA, USA, 2024, pp. 3891–3908.
8. E.M. Hutchins, M.J. Cloppert, R.M. Amin, “Intelligence-driven computer network defense informed by analysis of adversary campaigns and intrusion kill chains”, Proceedings of the 6th International Conference on Information Warfare and Security (ICIW), Washington, DC, USA, 2011, pp. 113–125.
9. B.D. Bryant, H. Saiedian, “A novel kill-chain framework for remote security log analysis with SIEM software”, Computers & Security, 2017, 67, 198–210. https://doi.org/10.1016/j.cose.2017.03.003
10. Mandiant, “M-Trends 2025”, Google Cloud Blog, 23 April 2025. Accessed 10 November 2026. https://cloud.google.com/blog/topics/threat-intelligence/m-trends-2025/
11. Cybersecurity and Infrastructure Security Agency (CISA), “Binding Operational Directive 22-01: Reducing the Significant Risk of Known Exploited Vulnerabilities”, CISA: Washington, DC, USA,

3 November 2021. Accessed 10 November 2026. https://www.cisa.gov/sites/default/files/publications/Reducing_the_Significant_Risk_of_Known_Exploited_Vulnerabilities_211103.pdf
12. T.D. Wagner, K. Mahbub, E. Palomar, A.E. Abdallah, "Cyber threat intelligence sharing: Survey and research directions", Computers & Security, 2019, 87, 101589.
13. The MITRE Corporation, "ATT&CK Updates – July 2020 (v7)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-july-2020
14. The MITRE Corporation, "ATT&CK Updates – October 2020 (v8)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-october-2020
15. The MITRE Corporation, "ATT&CK Updates – April 2021 (v9)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-april-2021
16. The MITRE Corporation, "ATT&CK Updates – October 2021 (v10)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-october-2021
17. The MITRE Corporation, "ATT&CK Updates – April 2022 (v11)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-april-2022/
18. The MITRE Corporation, "ATT&CK Updates – October 2022 (v12)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-october-2022/
19. The MITRE Corporation, "ATT&CK Updates – April 2023 (v13)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-april-2023
20. The MITRE Corporation, "ATT&CK Updates – October 2023 (v14)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-october-2023/
21. The MITRE Corporation, "ATT&CK Updates – April 2024 (v15)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-april-2024/
22. The MITRE Corporation, "ATT&CK Updates – October 2024 (v16)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-october-2024/
23. The MITRE Corporation, "ATT&CK Updates – April 2025 (v17)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-april-2025/
24. The MITRE Corporation, "ATT&CK Updates – October 2025 (v18)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-october-2025/
25. The MITRE Corporation, "ATT&CK Updates – April 2026 (v19)", Accessed 10 November 2026. https://attack.mitre.org/resources/updates/updates-april-2026/
26. E.H. Spafford, "The Internet worm program: An analysis", ACM SIGCOMM Computer Communication Review, 1989, 19, 17–57.
27. D. Moore, C. Shannon, K.C. Claffy, "Code-Red: A case study on the spread and victims of an Internet worm", Proceedings of the 2nd ACM SIGCOMM Workshop on Internet Measurement (IMW), Marseille, France, 2002, pp. 273–284.
28. D. Moore, V. Paxson, S. Savage, C. Shannon, S. Staniford, N. Weaver, "Inside the Slammer worm", IEEE Security & Privacy, 2003, 1, 33–39. https://doi.org/10.1109/MSECP.2003.1219056
29. D.E. Denning, "An intrusion-detection model", IEEE Transactions on Software Engineering, 1987, SE-13, 222–232. https://doi.org/10.1109/TSE.1987.232894
30. H. Debar, M. Dacier, A. Wespi, "Towards a taxonomy of intrusion-detection systems", Computer Networks, 1999, 31, 805–822. https://doi.org/10.1016/S1389-1286(98)00017-6

31. L. Bilge, T. Dumitraş, “Before we knew it: An empirical study of zero-day attacks in the real world”, Proceedings of the 2012 ACM Conference on Computer and Communications Security (CCS), Raleigh, NC, USA, 2012, pp. 833–844. https://doi.org/10.1145/2382196.2382284
32. R. Langner, “Stuxnet: Dissecting a cyberwarfare weapon”, IEEE Security & Privacy, 2011, 9, 49–51. https://doi.org/10.1109/MSP.2011.67
33. A. Kharraz, W. Robertson, D. Balzarotti, L. Bilge, E. Kirda, “Cutting the Gordian knot: A look under the hood of ransomware attacks”, Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA 2015), Lecture Notes in Computer Science, Volume 9148, Springer: Cham, Switzerland, 2015, pp. 3–24. https://doi.org/10.1007/978-3-319-20550-2_1
34. F. Barr-Smith, X. Ugarte-Pedrero, M. Graziano, R. Spolaor, I. Martinovic, “Survivalism: Systematic analysis of Windows malware living-off-the-land”, Proceedings of the 2021 IEEE Symposium on Security and Privacy (SP), 2021, pp. 1557–1574. https://doi.org/10.1109/SP40001.2021.00047
35. CISA, NSA, FBI, et al., “Identifying and Mitigating Living Off the Land Techniques”, Joint Guidance, February 2024. Accessed 10 November 2026. https://www.cisa.gov/resources-tools/resources/identifying-and-mitigating-living-land-techniques
36. Cybersecurity and Infrastructure Security Agency (CISA), “Emergency Directive 21-01: Mitigate SolarWinds Orion Code Compromise”, 13 December 2020. Accessed 10 November 2026. https://www.cisa.gov/emergency-directive-21-01
37. M. Ohm, H. Plate, A. Sykosch, M. Meier, “Backstabber's knife collection: A review of open source software supply chain attacks”, arXiv, 2020, arXiv:2005.09535. https://arxiv.org/abs/2005.09535
38. H. Xu, S. Wang, N. Li, K. Wang, Y. Zhao, K. Chen, T. Yu, Y. Liu, H. Wang, “Large language models for cyber security: A systematic literature review”, ACM Transactions on Software Engineering and Methodology, 2025, arXiv:2405.04760. https://arxiv.org/abs/2405.04760
39. I. Hasanov, S. Virtanen, A. Hakkala, J. Isoaho, “Application of large language models in cybersecurity: A systematic literature review”, IEEE Access, 2024, 12, 176751–176778. https://doi.org/10.1109/ACCESS.2024.3505983
40. R. Fang, R. Bindu, A. Gupta, D. Kang, “LLM agents can autonomously exploit one-day vulnerabilities”, arXiv, 2024, arXiv:2404.08144. https://arxiv.org/abs/2404.08144
41. Anthropic, “Disrupting the First Reported AI-Orchestrated Cyber Espionage Campaign”, 13 November 2025. Accessed 10 November 2026. https://www.anthropic.com/news/disrupting-AI-espionage
42. H. Pearce, B. Ahmad, B. Tan, B. Dolan-Gavitt, R. Karri, “Asleep at the keyboard? Assessing the security of GitHub Copilot's code contributions”, Proceedings of the 2022 IEEE Symposium on Security and Privacy (SP), San Francisco, CA, USA, 2022, pp. 754–768. https://doi.org/10.1109/SP46214.2022.9833571
43. J. Spracklen, R. Wijewickrama, A.H.M.N. Sakib, A. Maiti, B. Viswanath, M. Jadliwala, “We have a package for you! A comprehensive analysis of package hallucinations by code generating LLMs”, Proceedings of the 34th USENIX Security Symposium, Seattle, WA, USA, 2025, pp. 3687–3706.
44. S. Bhatt, P.K. Manadhata, L. Zomlot, “The operational role of security information and event management systems”, IEEE Security & Privacy, 2014, 12, 35–41. https://doi.org/10.1109/MSP.2014.103
45. W. Tounsi, H. Rais, “A survey on technical threat intelligence in the age of sophisticated cyber attacks”, Computers & Security, 2018, 72, 212–233. https://doi.org/10.1016/j.cose.2017.09.001

46. V. Chandola, A. Banerjee, V. Kumar, “Anomaly detection: A survey”, ACM Computing Surveys, 2009, 41, 15. https://doi.org/10.1145/1541880.1541882
47. R. Sommer, V. Paxson, “Outside the closed world: On using machine learning for network intrusion detection”, Proceedings of the 2010 IEEE Symposium on Security and Privacy, 2010, pp. 305–316.
48. SigmaHQ, “Sigma: Generic Signature Format for SIEM Systems”, Accessed 10 November 2026. https://github.com/SigmaHQ/sigma
49. M. Vermeer, M. van Eeten, C. Gañán, “Ruling the rules: Quantifying the evolution of rulesets, alerts and incidents in network intrusion detection”, Proceedings of the 2022 ACM Asia Conference on Computer and Communications Security (ASIA CCS), 2022, pp. 799–813. https://doi.org/10.1145/3488932.3517412
50. K.T.W. Teuwen, T. Mulders, E. Zambon, L. Allodi, “Ruling the unruly: Designing effective, low-noise network intrusion detection rules for security operations centers”, Proceedings of the 20th ACM Asia Conference on Computer and Communications Security (ASIA CCS), Hanoi, Vietnam, 2025, pp. 1428–1441. https://doi.org/10.1145/3708821.3710823
51. F. Valeur, G. Vigna, C. Kruegel, R.A. Kemmerer, “A comprehensive approach to intrusion detection alert correlation”, IEEE Transactions on Dependable and Secure Computing, 2004, 1, 146–169. https://doi.org/10.1109/TDSC.2004.21
52. P. Ning, Y. Cui, D.S. Reeves, “Constructing attack scenarios through correlation of intrusion alerts”, Proceedings of the 9th ACM Conference on Computer and Communications Security (CCS), Washington, DC, USA, 2002, pp. 245–254. https://doi.org/10.1145/586110.586144
53. S.M. Milajerdi, R. Gjomemo, B. Eshete, R. Sekar, V.N. Venkatakrishnan, “HOLMES: Real-time APT detection through correlation of suspicious information flows”, Proceedings of the 2019 IEEE Symposium on Security and Privacy (SP), 2019, pp. 1137–1152. https://doi.org/10.1109/SP.2019.00026
54. W.U. Hassan, S. Guo, D. Li, Z. Chen, K. Jee, Z. Li, A. Bates, “NoDoze: Combatting threat alert fatigue with automated provenance triage”, Proceedings of the Network and Distributed System Security Symposium (NDSS), San Diego, CA, USA, 2019. https://doi.org/10.14722/ndss.2019.23349
55. S. Axelsson, “The base-rate fallacy and the difficulty of intrusion detection”, ACM Transactions on Information and System Security, 2000, 3, 186–205. https://doi.org/10.1145/357830.357849
56. L. Yang, Z. Chen, C. Wang, Z. Zhang, S. Booma, P. Cao, C. Adam, A. Withers, Z. Kalbarczyk, R.K. Iyer, G. Wang, “True attacks, attack attempts, or benign triggers? An empirical measurement of network alerts in a security operations center”, Proceedings of the 33rd USENIX Security Symposium, Philadelphia, PA, USA, 2024, pp. 1525–1542.
57. B.A. Alahmadi, L. Axon, I. Martinovic, “99% false positives: A qualitative study of SOC analysts' perspectives on security alarms”, Proceedings of the 31st USENIX Security Symposium, Boston, MA, USA, 2022, pp. 2783–2800.
58. F.B. Kokulu, A. Soneji, T. Bao, Y. Shoshitaishvili, Z. Zhao, A. Doupé, G.-J. Ahn, “Matched and mismatched SOCs: A qualitative study on security operations center issues”, Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (CCS), London, UK, 2019, pp. 1955–1970. https://doi.org/10.1145/3319535.3354239
59. S.C. Sundaramurthy, A.G. Bardas, J. Case, X. Ou, M. Wesch, J. McHugh, S.R. Rajagopalan, “A human capital model for mitigating security analyst burnout”, Proceedings of the Eleventh

Symposium on Usable Privacy and Security (SOUPS 2015), Ottawa, ON, Canada, 2015, pp. 347–359.

60. T.-F. Yen, A. Oprea, K. Onarlioglu, T. Leetham, W. Robertson, A. Juels, E. Kirda, “Beehive: Large-scale log analysis for detecting suspicious activity in enterprise networks”, Proceedings of the 29th Annual Computer Security Applications Conference (ACSAC), New Orleans, LA, USA, 2013, pp. 199–208. https://doi.org/10.1145/2523649.2523670
61. G. Ho, M. Dhiman, D. Akhawe, V. Paxson, S. Savage, G.M. Voelker, D. Wagner, “Hopper: Modeling and detecting lateral movement”, Proceedings of the 30th USENIX Security Symposium, 2021, pp. 3093–3110.
62. R. Singh, S. Tariq, F. Jalalvand, M.B. Chhetri, S. Nepal, C. Paris, M. Lochner, “LLMs in the SOC: An empirical study of human-AI collaboration in security operations centres”, arXiv, 2025, arXiv:2508.18947. https://arxiv.org/abs/2508.18947
63. S.A. Tsegai, A. Kantchelian, D. Yao, P. Gao, “From investigation failures to reliable SOC agents: Understanding and improving LLM-based alert triage”, arXiv, 2026, arXiv:2610.10608. https://arxiv.org/abs/2610.10608