Open Access Library Journal
2026, Volume 13, e16049
ISSN Online: 2333-9721
ISSN Print: 2333-9705

Beyond Authentication: A Systematic Review and Identity-Centric MITRE ATT&CK Framework for Detecting Malicious Post-Authentication Behaviour

Samuel Dushimimana Shyaka[1], Collin Arnold Kabwama[2], Dushime Benjamin[3]

[1]College of Engineering, Carnegie Mellon University Africa, Kigali, Rwanda
[2]Department of Computer Science, Maharishi International University, Fairfield, IA, USA
[3]Raymond A. Mason School of Business, William & Mary, Williamsburg, VA, USA
Email: sdushimi@alumni.cmu.edu

How to cite this paper: Shyaka, S.D., Kabwama, C.A. and Benjamin, D. (2026) Beyond Authentication: A Systematic Review and Identity-Centric MITRE ATT&CK Framework for Detecting Malicious Post-Authentication Behaviour. *Open Access Library Journal*, **13**: e16049.
https://doi.org/10.4236/oalib.1116049

Received: September 20, 2026
Accepted: November 17, 2026
Published: November 20, 2026

Copyright © 2026 by author(s) and Open Access Library Inc.
This work is licensed under the Creative Commons Attribution International License (CC BY 4.0).
http://creativecommons.org/licenses/by/4.0/



Open Access

Abstract

Credential compromise lets an adversary operate through accounts, tokens, and sessions that appear legitimate. This article reports a PRISMA 2020 systematic review of research on detecting malicious post-authentication behaviour and proposes an identity-centric framework grounded in MITRE ATT&CK. Six predeclared Semantic Scholar queries (2010 to September 2026) and citation searching yielded 135 included studies. A six-domain methodological appraisal was applied to the 37 studies evaluated on public benchmarks. The evidence shows rapid growth (57 studies appeared in 2025 or 2026), a shift toward graph and hybrid deep models, and persistent weaknesses: among the 119 evaluative detection studies, chronological validation is reported by 13 percent, false-positive rates or alert budgets by 28 percent, and ATT&CK use or attack-sequence modelling by 14 percent. Graph, sequence, and tensor models evaluated on the Los Alamos National Laboratory (LANL) datasets, whose labels come from a red-team exercise, provide the most credible evidence available on detection at low false-positive rates; the review does not directly compare cross-source with authentication-only detection, which remains a hypothesis for the proposed protocol. Building on these findings, the article defines malicious post-authentication behaviour operationally, maps identity-related ATT&CK pathways to observable evidence, specifies a detection architecture with a calibrated and explainable session-risk formulation, and sets out a reproducible validation protocol for the LANL multi-source dataset.

Subject Areas

Computer Science, Information Security

Keywords

Valid Accounts, Post-Authentication Detection, MITRE ATT&CK, Zero Trust, User and Entity Behaviour Analytics, Lateral Movement, Compromised Account Detection, Systematic Review

1. Introduction

Identity has become a principal control plane for enterprise and cloud security. Credential abuse remained the most common initial-access vector in the 2025 Verizon Data Breach Investigations Report, appearing in 22 percent of analysed breaches [1], and Microsoft reported that identity-based attacks rose by 32 percent in the first half of 2025, with more than 97 percent of them being password attacks [2]. These figures describe pressure at the authentication boundary. The harder defensive problem begins after an attacker holds a valid password, token, certificate, session cookie, service credential, or federated assertion, because every subsequent action is performed through an identity that the access-control system has already accepted.

MITRE ATT&CK lists Valid Accounts (T1078) under Initial Access, Persistence, Privilege Escalation, and Defense Evasion, so the same identity can be both the entry mechanism and the vehicle for later actions [3], and attackers dwell in hijacked enterprise accounts for days to weeks [4]. NIST SP 800-207 states that no implicit trust should be granted on the basis of network location or ownership [5], and the CISA Zero Trust Maturity Model places identity alongside devices, networks, applications, and data [6]. Yet many detection studies analyse an abnormal login time, device, or destination without connecting the login to the ordered behaviours that follow it, and fewer explain those behaviours with an adversary knowledge base.

Existing reviews cover parts of this problem: insider-threat detection [7] [8], deep learning for log anomaly detection [9], ATT&CK research [10], and compromised accounts in online social networks [11] [12]. None treats the authenticated session as the unit of analysis or examines how studies validate claims under the severe class imbalance and temporal dependence of identity telemetry. This article addresses that gap through a PRISMA 2020 systematic review [13] and an implementable framework. It contributes a reproducible evidence base of 135 studies; a thematic synthesis with a methodological appraisal of the benchmark-evaluated literature; an operational definition of malicious post-authentication behaviour with an identity-centric ATT&CK mapping; a detection architecture with a calibrated, explainable session-risk formulation; and a validation protocol for public datasets.

2. Systematic Review Method

2.1. Design, Research Questions, and Search Strategy

The review followed the PRISMA 2020 structure [13] and the guidance of Kitch-

enham *et al.* for software-engineering reviews [14]. Four research questions were addressed: which identity and post-authentication threat models are represented (RQ1); which telemetry, features, and analytical methods support detection (RQ2); how studies validate performance under class imbalance and temporal dependence (RQ3); and how ATT&CK and zero-trust principles can support interpretable, risk-adaptive response (RQ4).

Records were retrieved from the Semantic Scholar Academic Graph bulk search interface, which supports re-executable Boolean queries over titles and abstracts. Searches were run on 6 September 2026 and limited to computer-science records published between 1 January 2010 and 2 September 2026. Six predeclared query concepts are listed in **Table 1**; the full search log is provided in supplementary table S1 (sheets S1a and S1b of the supplementary workbook submitted with this article as a separate file). Backward citation tracing from recent reviews and from studies using the Los Alamos National Laboratory (LANL) datasets identified studies missed by the database search. The computer-science filter shapes coverage: work indexed only under engineering, management, information systems, or sector-specific fields (energy, water, transport, health) was not retrieved, so cloud-governance, organisational, and critical-infrastructure research is likely underrepresented. Only 6 of the 135 included studies address cloud platforms and none is coded as a critical-infrastructure study, so Section 10 draws implications for that setting without reviewing its literature.

Table 1. Predeclared search concepts, query strings, and yields.

Concept	Query string (Semantic Scholar bulk search syntax)	Records	Included
Valid accounts	"valid account" | "valid accounts"	27	2
Authentication anomaly detection	"authentication anomaly detection" | "authentication log" + anomaly | "login anomaly"	56	17
User and entity behaviour analytics	"user and entity behavior analytics" | "user entity behavior analytics" | UEBA | "user behavior analytics" + security	153	35
Compromised accounts	"compromised account detection" | "compromised account" + detection | "account takeover detection" | "malicious login"	153	49
Lateral movement	"lateral movement detection" | "lateral movement" + authentication + detection	72	29
ATT&CK and identity	"MITRE ATT&CK" + ("valid accounts" | credential | identity) + detection	23	5
Total (before deduplication)		484	130 (unique)

Included counts are attributed to every query that returned the record, so they sum to more than the number of unique studies. Filters: field of study = Computer Science; years 2010 to 2026.

2.2. Eligibility, Screening, and Extraction

A record was eligible when it (i) was a peer-reviewed journal, conference, or workshop paper; (ii) addressed malicious or anomalous use of an already authenticated or legitimately credentialed identity, including hijacked accounts, malicious

logins, credential misuse, insider anomalies, credential-based lateral movement, user and entity behaviour analytics (UEBA), session-level analysis, or ATT&CK-mapped detection of identity techniques; (iii) presented an evaluable contribution (a detection method with results, a dataset, a framework with a worked evaluation, or a review); and (iv) concerned enterprise, organisational, or cloud IT, or an online-service account-takeover setting whose method analyses behaviour within an authenticated session. Excluded were biometric continuous authentication without a threat-detection component, IoT or physical-layer device authentication, fake or bot account detection, generic intrusion, malware, or phishing studies without an identity focus, and non-technical items.

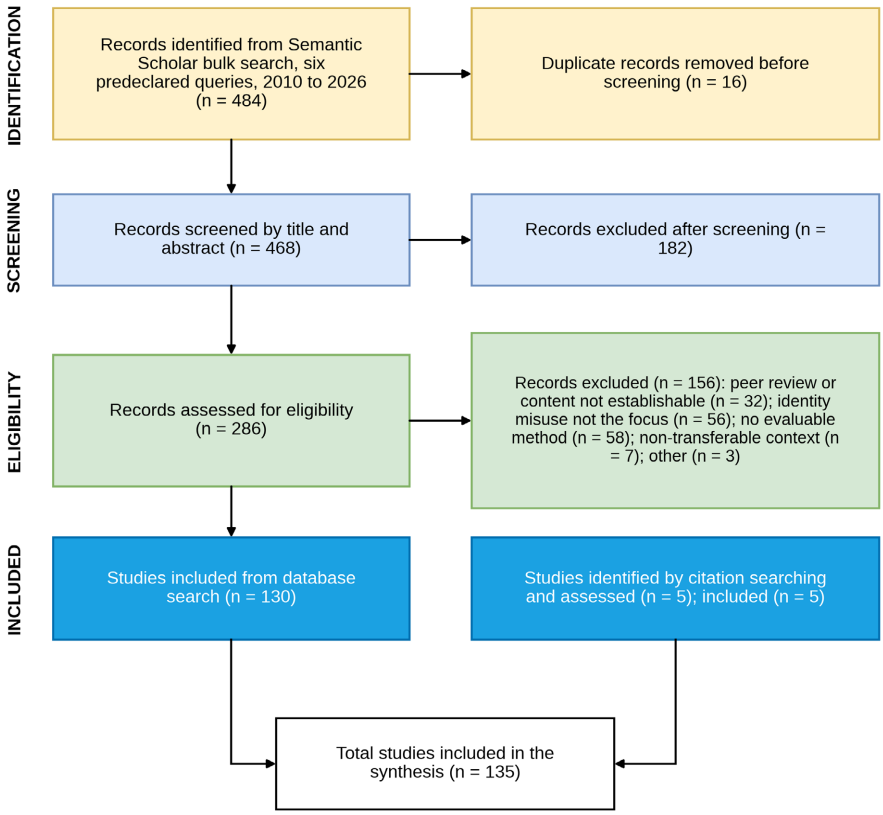


Figure 1. PRISMA 2020 flow of identification, screening, eligibility assessment, and inclusion.

Records were deduplicated by DOI and normalised title, screened on title and abstract, and assessed for eligibility with a recorded reason code for each exclusion (supplementary table S1); rule-assisted first-pass classification was reviewed decision by decision by the author, and 56 records without an abstract were assessed from the title and, where known, the full text (**Figure 1**). A record assessed on title alone was retained only when the title itself established the identity-misuse focus and an evaluable contribution; otherwise it was excluded under the code for content that could not be established (32 records). For records with an abstract, eligibility was decided on the abstract, and the full text was consulted for the studies

taken forward to detailed synthesis and appraisal. Eligibility was assessed from the full text or known content for 31 of the 130 database-retrieved studies and for the five citation-searched studies, 36 of 135 in all, as recorded in the screening record; the other 99 were judged eligible on title and abstract, so their eligibility and coded characteristics rest on what the abstract reports. Extracted items were year, context, identity type, threat model, telemetry, method family, dataset, validation design, class-imbalance treatment, false-positive or alert-budget reporting, metrics, ATT&CK use, principal finding, and limitations; the extraction for every included study is given in supplementary table S2. Heterogeneity precluded statistical pooling, so the synthesis combined descriptive mapping with thematic grouping.

2.3. Methodological Appraisal

Because conventional risk-of-bias instruments do not fit security detection studies, a six-domain appraisal was applied to the 37 studies evaluated on a public benchmark with documented ground truth (LANL, CERT insider-threat, DARPA Transparent Computing and OpTC, or comparable sources), for which the appraisal can be checked by others. Fifty included studies used a public data source. The 37 appraised studies are all 35 that used LANL, CERT, or DARPA data (22, 12, and 1) and the two using another public source whose labels are documented. The other 13 were not appraised: three are reviews or mappings, and ten use sources whose label provenance and class balance (domains D2 and D3) could not be checked. The two LANL dataset descriptions [15] [16] are scored for completeness. Each domain was scored 0 (not addressed), 1 (partially), or 2 (clearly addressed): D1 chronological separation of training and evaluation windows; D2 label provenance from red-team activity or confirmed incidents rather than synthetic injection; D3 realistic class rarity in the evaluation set; D4 external validation on more than one dataset, site, or period; D5 reporting of false-positive rate, precision at a fixed alert budget, alerts per analyst-day, or time to detection; and D6 reproducibility through public data or code. Totals of 10 to 12 were rated high, 6 to 9 moderate, and 0 to 5 limited. Domain scores and a justification for each appraised study are given in supplementary table S3. The score indicates evidence maturity, not scientific merit.

3. Results: Search Yield and Study Characteristics

The six queries returned 484 records; after deduplication, screening, and eligibility assessment (exclusion reasons in **Figure 1**), 130 studies remained. Citation searching added five: Eren *et al.* [17], King and Huang [18], Log2vec [19], and the two LANL dataset descriptions [15] [16], giving 135 included studies.

Publication activity has accelerated sharply: 57 of the 135 studies (42 percent) appeared in 2025 or 2026, compared with 33 in 2012 to 2019 (**Figure 2**). Eighty-eight studies addressed enterprise environments, 27 online services, 6 cloud platforms, 12 were reviews, and 2 covered mixed settings. Context and method family were coded separately, which explains the two review counts: 12 studies are re-

views with no single deployment context, while 16 fall in the review method family because that family also holds four studies coded to a context (the two LANL dataset descriptions [15] [16], an ATT&CK identity-threat mapping [20], and one mixed-setting overview). The remaining 119 are evaluative detection studies. Human identities were the sole subject of 71 studies, 56 covered both human and machine principals (typically through user-to-computer authentication graphs), and only one concentrated on service or machine identities; identity type was not reported for the remaining seven. By threat model, 35 studies targeted compromised accounts, 28 lateral movement, 26 insider misuse, 22 malicious logins, 18 general UEBA, and 6 account takeover. Statistical and rule-based approaches (30) and classical machine learning (29) remain common, but ensemble and hybrid architectures (26) and graph learning (19) dominate the most recent period; deep sequence models (9), other deep models (3), and tensor factorisation (3) form smaller groups, and 16 studies were reviewsor dataset descriptions, so the eight families sum to 135.

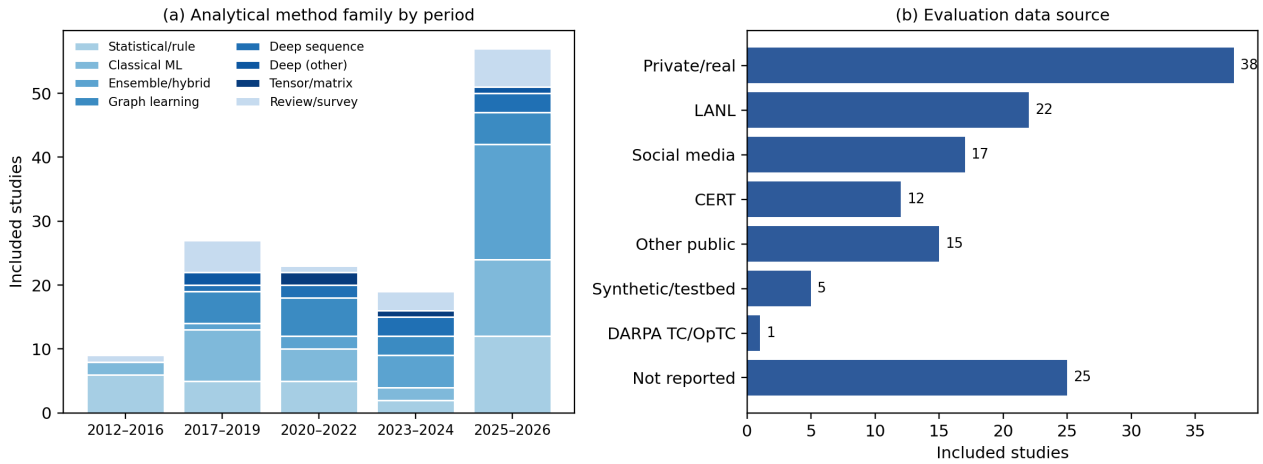


Figure 2. Evidence map of the 135 included studies: (a) analytical method family by publication period; (b) evaluation data source.

Evaluation data matter most. Thirty-eight studies used private data, 25 did not report a dataset, 22 used a LANL dataset, 17 used self-collected social-media data, 12 a CERT insider-threat release, one DARPA Transparent Computing or OpTC data, 15 other public sources, and 5 synthetic or testbed data. LANL and CERT are thus the only substantial bodies of comparable work, and they differ in a way that matters: LANL red-team labels record authentication events generated by an authorised red-team exercise on the operational network, not naturally occurring adversary activity, whereas CERT labels are synthetic scenario injections [21]. Red-team labels are therefore more realistic than injected scenarios but still reflect the tactics, timing, and scope chosen by one exercise team.

4. Thematic Synthesis of Detection Research

4.1. Threat Models Represented in the Literature (RQ1)

The literature approaches the post-authentication problem from four directions

that are rarely connected. The first is the compromised or hijacked account, studied mostly in online services. Per-user behavioural, clickstream, and posting profiles detect compromised social accounts at low false-positive rates [22]. Later work added change-point detection over posting behaviour [23].

The second direction is the malicious login in an enterprise authentication graph, most directly relevant to valid-account abuse. Siadati and Memon defined structurally anomalous logins in terms of user, host, and role attributes [24]. Kaiafas *et al.* trained supervised ensembles on LANL authentication events with attention to trustworthy false-positive reporting [25]. Social-Hunter combines social heuristics with authentication context to unveil unknown malicious logins using valid accounts at low false-positive rates on the LANL multi-source data [26].

The third direction is lateral movement, which treats the login as one edge in an evolving user-to-computer graph. Bowman *et al.* detected lateral movement without labels using unsupervised graph learning over LANL authentication graphs [27]; Bian *et al.* engineered graph features with attention to sampling and overhead [28]; and Euler scaled temporal link prediction to full LANL and OpTC data with substantially faster training than prior temporal graph models [18].

The fourth direction is the malicious or careless insider, studied almost entirely on the CERT synthetic datasets [21]. Deep, federated, gradient-boosting, and language-model approaches report high accuracy on these releases [29]-[32]. A compromised account and a malicious insider produce similar telemetry, but the label semantics differ (Section 5). Two smaller groups complete the picture: UEBA systems that score behaviour against personal and peer baselines [33], and studies that connect detection with zero-trust response through continuous risk scoring that triggers step-up challenges or session termination [34].

4.2. Telemetry, Features, and Analytical Methods (RQ2)

Statistical, peer-group, and tensor models. A login to a new host can be benign as a point event yet suspicious when it follows token issuance from an unusual device and precedes rapid access to many servers. Peer grouping reduces the cold-start problem and materially reduces first-access false positives [33], but can also reproduce organisational drift or mask coordinated misuse. On the LANL datasets [15] [16], Goodman *et al.* derived bipartite anomaly features from user-computer graphs [35], and Eren *et al.* showed that non-negative tensor factorisation detects red-team authentication and process anomalies without labels, reporting precision at low false-positive rates on chronologically separated data [17]. Simple rule-based detectors remain competitive for high-rate attacks, while learned models add value for stealthy activity [36].

Sequence and log-language models. Log-language models learn event ordering, although performance remains sensitive to parsing, windowing, and distribution stability [9]. Gökstorp *et al.* trained a self-supervised transformer on LANL authentication events with a chronological split and reported an average precision

of 0.041 at an AUC above 0.99, illustrating how uninformative AUC is under extreme rarity [37].

Graph-based detection. Authentication forms a temporal heterogeneous graph of users, devices, processes, and resources, and graph methods capture lateral movement and rare relationships that flat feature vectors obscure; the same advantage is reported in fraud detection [38] [39]. Log2vec reported high AUC on CERT and LANL data [19], Euler achieved strong precision on LANL and OpTC [18], and newer models add temporal attention, drift-triggered retraining, and provenance graphs across hybrid cloud telemetry [40] [41]. Privilege-transition graphs rooted at identities outperform feature-vector baselines on LANL data, yet their precision-recall AUC of 0.331 shows how far the problem is from solved at realistic base rates [42]. Enterprise graphs are dynamic and partially observed, and evaluation can leak future topology unless splits are strictly temporal [18].

Hybrid, ensemble, and zero-trust integrated systems. Recent studies combine transformer and graph ensembles that cut false positives relative to rule-based SIEM [43], apply dual-granularity isolation forests to 403 million GitHub audit records [44], and re-score session behaviour on device to trigger step-up challenges [34]. These systems are closest to the framework in Section 7, but most, like emerging language-model detectors [32] [45], are evaluated on private or synthetic data and rarely report performance under a fixed alert budget.

4.3. ATT&CK Integration

ATT&CK is used mainly for threat intelligence, hunting, modelling, and simulation [10], and only 19 included studies use ATT&CK or explicit attack-sequence modelling. Serzantov *et al.* mapped 366 identity-related ATT&CK threats across the digital identity lifecycle [20]; Zhang *et al.* combined Sigma rules, ATT&CK enrichment, and graph analysis to reconstruct attack chains from LANL data [46]; and several lateral-movement detectors present ATT&CK-mapped explanations [41]. Mapping low-level logs to techniques remains uncertain and individual events support several interpretations, so an identity-centric design must report the evidentiary strength and uncertainty of each mapping.

5. Validation Practices and Methodological Appraisal (RQ3)

5.1. Reporting Practices across the Corpus

Reporting practices are expressed against the 119 evaluative detection studies, excluding the 16 reviews and dataset descriptions, because validation practices do not apply to the latter. At title-and-abstract level, chronological separation of training and evaluation data was reported by 15 of 119 studies (13 percent) and false-positive rates, precision at a fixed alert budget, or comparable operational quantities by 33 (28 percent). Seventeen (14 percent) used ATT&CK or attack-sequence modelling, a measure of use that does not imply an ATT&CK-aligned explanation of alerts, and 45 (38 percent) used a public benchmark. Against all 135 included studies the corresponding figures are 15 (11 percent), 34 (25 per-

cent), 19 (14 percent), and 50 (37 percent), as plotted in Figure 3. Explicit class-imbalance treatment beyond resampling was reported by 10 evaluative studies (8 percent) and accuracy by 28 (24 percent); accuracy is close to meaningless when the positive class is below one event per ten thousand. These are lower bounds because abstracts omit detail, but they are consistent with long-standing warnings about the base-rate fallacy [47].

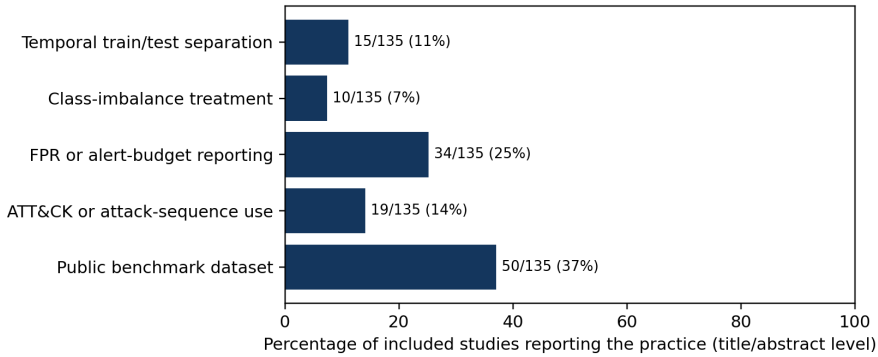


Figure 3. Reporting of validation practices across the 135 included studies (title-and-abstract level); percentages use all included studies as the denominator, and values for the 119 evaluative studies are given in the text.

5.2. Appraisal of Benchmark-Evaluated Studies

Table 2 reports the appraisal of the 37 benchmark-evaluated studies. Three scored in the high band: Bowman *et al.* [27], King and Huang [18], and Eren *et al.* [17], all combining chronological evaluation, red-team labels, realistic rarity, multi-dataset validation, alert-level reporting, and public data or code. Ten scored moderate and 24 limited (Table 2). Chronological separation had the lowest mean score (0.43 of 2), followed by class-balance realism (0.51) and external validation (0.68); label provenance scored highest (1.49) because LANL red-team labels are widely used; domain D2 credits red-team exercises as documented ground truth, which is weaker evidence than confirmed real-world incidents. Almost all CERT-based insider studies scored limited, not because the models are weak but because CERT labels are synthetic and most studies balance the test set, so their reported accuracies of 96 to 99 percent do not estimate operational performance. LANL-based studies that report precision at a low false-positive rate consistently show precision-recall AUC values well below 0.5 even for strong models [37] [42]. Claims of superiority resting on accuracy or AUC alone should therefore be discounted.

Table 2. Methodological appraisal of the 37 benchmark-evaluated studies (D1 chronological separation, D2 label provenance, D3 class-balance realism, D4 external validation, D5 operational alert reporting, D6 reproducibility; 0 to 2 each; OpTC = DARPA Transparent Computing/OpTC).

Study	Year	Data	D1	D2	D3	D4	D5	D6	Sum	Band	Study	Year	Data	D1	D2	D3	D4	D5	D6	Sum	Band
Goodman *et al.* [35]	2015	LANL	2	2	1	0	2	1	8	Moderate	Artioli *et al.* [48]	2024	CERT	0	1	0	2	0	1	4	Limited
Kent [15]	2016	LANL	0	2	0	0	0	2	4	Limited	Gökstorp *et al.* [37]	2024	LANL	2	2	2	0	1	1	8	Moderate

Continued

Siadati *et al.* [49]	2016	LANL	0	2	0	0	0	1	3	Limited	Aljumaily and Abd [43]	2025	CERT	0	1	0	2	1	2	6	Moderate
Kaiafas *et al.* [25]	2018	LANL	2	2	1	0	2	1	8	Moderate	Kodumuri and Dadi [50]	2025	CERT	0	1	0	0	0	1	2	Limited
Turcotte *et al.* [16]	2018	LANL	0	0	0	0	0	2	2	Limited	Mohideen *et al.* [51]	2025	Other	0	1	0	0	0	1	2	Limited
Amrouche *et al.* [52]	2019	LANL	0	2	0	0	0	1	3	Limited	Sunny *et al.* [53]	2025	CERT	0	1	0	0	0	1	2	Limited
Liu *et al.* [19]	2019	LANL	0	2	1	2	2	2	9	Moderate	Tang and Ge [26]	2025	LANL	2	2	2	0	2	1	9	Moderate
Bowman *et al.* [27]	2020	LANL	2	2	2	2	2	2	12	High	Zhang *et al.* [46]	2025	LANL	0	2	0	0	0	1	3	Limited
Eren *et al.* [54]	2020	LANL	0	2	1	0	0	1	4	Limited	Cheng *et al.* [32]	2026	CERT	0	1	0	0	1	1	3	Limited
Kaiafas *et al.* [55]	2020	LANL	0	2	1	0	2	1	6	Moderate	Thakare *et al.* [56]	2026	CERT	0	1	0	2	1	1	5	Limited
Sharma *et al.* [29]	2020	CERT	0	1	0	0	2	1	4	Limited	Goud *et al.* [42]	2026	LANL	0	2	2	0	2	1	7	Moderate
Wei *et al.* [57]	2020	LANL	0	2	0	0	0	1	3	Limited	Maltsev & Budovskaya [58]	2026	CERT	0	1	0	0	0	1	2	Limited
Zhao *et al.* [59]	2020	LANL	2	2	0	0	0	1	5	Limited	Mozumder *et al.* [60]	2026	CERT	0	1	0	2	0	1	4	Limited
Bian *et al.* [28]	2021	LANL	0	2	0	2	0	1	5	Limited	Muratović and Isaković [36]	2026	LANL	0	1	2	0	2	1	6	Moderate
Wu *et al.* [61]	2021	LANL	0	2	0	0	0	1	3	Limited	Nina *et al.* [41]	2026	OpTC	0	2	0	2	2	1	7	Moderate
Eren *et al.* [17]	2023	LANL	2	2	2	2	2	2	12	High	Athapaththu *et al.* [62]	2026	LANL	0	1	0	2	1	1	5	Limited
King and Huang [18]	2023	LANL	2	2	2	2	2	2	12	High	Paduvilan *et al.* [30]	2026	CERT	0	1	0	0	2	1	4	Limited
											Samuel *et al.* [31]	2026	CERT	0	1	0	2	1	1	5	Limited
											Thammaiah *et al.* [63]	2026	Other	0	0	0	0	1	1	2	Limited
											Trinita & Kurniawan [64]	2026	CERT	0	1	0	1	1	1	4	Limited

6. Conceptual Foundations and Identity-Related ATT&CK Pathways

Authentication verifies evidence associated with an identity at session establishment or reauthentication; authorisation determines whether that identity may perform a requested action; post-authentication detection evaluates whether the sequence of authenticated and authorised actions remains consistent with legitimate intent, role, device posture, resource sensitivity, and policy. Multifactor authentication reduces account takeover, yet cookie theft, token replay, compromised endpoints, and insiders bypass or follow successful authentication.

Malicious post-authentication behaviour is therefore defined as a sequence of actions performed through a successfully authenticated human or non-human identity that advances an adversary objective or violates the identity's legitimate task intent. The definition includes abuse of valid credentials, hijacked sessions, delegated tokens, dormant accounts, service principals, and malicious insiders, and excludes a failed sign-in unless it supplies context for a later successful session. **Figure 4** shows the resulting detection boundary.

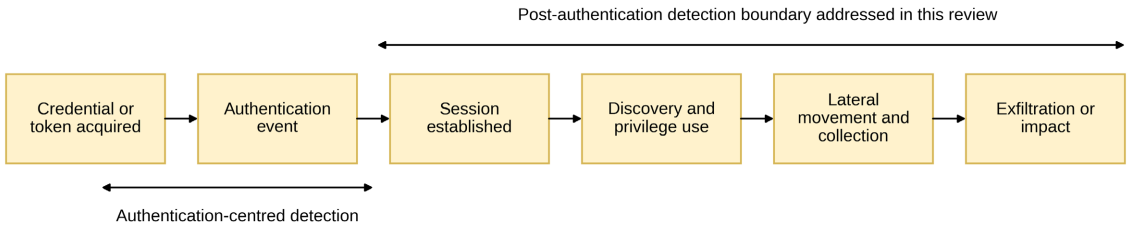


Figure 4. Detection boundary from credential acquisition through continuous evaluation of the authenticated session.

T1078 is the analytical anchor within a broader pathway [3]: credential access may precede valid-account use, followed by account and permission discovery, remote-service activity, collection, and exfiltration. Detection therefore requires path inference from incomplete, heterogeneous telemetry, a problem made concrete for identity threats by Serzantov *et al.* [20]. **Table 3** organises the pathway by stage.

Table 3. Identity-centric ATT&CK pathway and required evidence.

Stage	Representative ATT&CK behaviours	Observable evidence	Interpretive value
Acquire identity material	Brute Force (T1110); Credentials from Password Stores (T1555); Steal Web Session Cookie (T1539)	Failed and successful logons, token issuance, cookie use, credential-store access	Explains how authentication material became adversary-controlled
Establish session	Valid Accounts (T1078); External Remote Services (T1133)	Identity-provider event, source and device, token claims, VPN or cloud session	Defines the authenticated identity and initial context
Orient and expand	Account Discovery (T1087); Permission Groups Discovery (T1069); System Network Connections Discovery (T1049)	Directory queries, command-line and process events, API calls, enumeration volume	Signals reconnaissance after access
Move and escalate	Remote Services (T1021); Account Manipulation (T1098); Access Token Manipulation (T1134)	New destinations, remote logons, group changes, token changes	Links identity use to propagation or privilege gain
Collect and remove	Data from Local System (T1005); Archive Collected Data (T1560); Exfiltration Over Web Service (T1567)	File reads, compression, unusual egress, cloud uploads	Connects account activity to mission impact
Disrupt or conceal	Impair Defenses (T1562); Inhibit System Recovery (T1490); Account Access Removal (T1531)	Logging changes, recovery changes, account disablement	Raises severity and informs rapid containment

7. Proposed Identity-Centric Detection Framework

The framework (**Figure 5**) synthesises the strongest elements of the reviewed literature and is presented as a research design to be tested, not as a validated system.

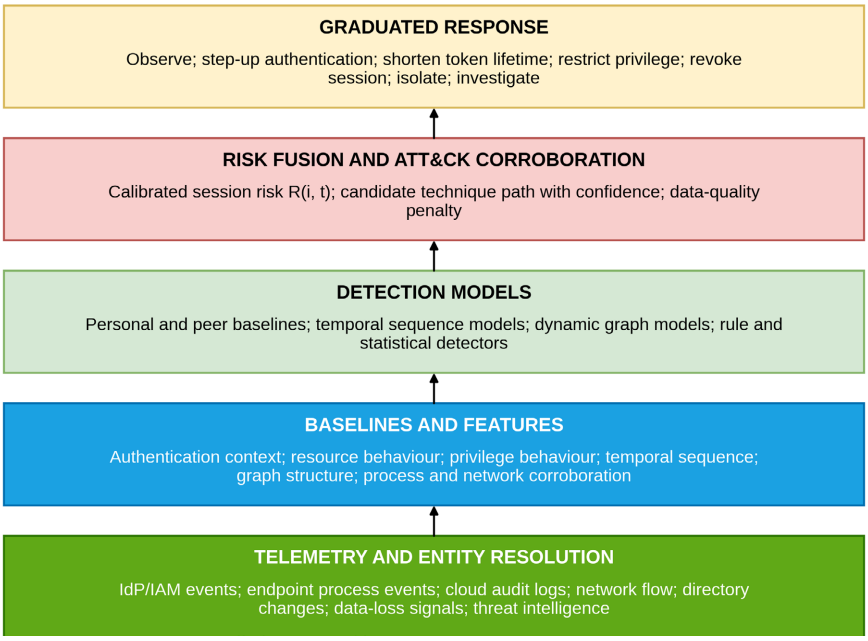


Figure 5. Layered architecture for ATT&CK-informed post-authentication detection.

7.1. Telemetry, Entity Resolution, and Features

The telemetry layer combines identity-provider and IAM events, endpoint process events, cloud audit logs, network flow, directory changes, application access, data-loss signals, and threat intelligence. A canonical event schema should record time, actor, credential or token, source, destination, action, outcome, privilege context, device posture, session identifier, and provenance, and entity resolution must preserve links among accounts, federated identities, service principals, and devices [16]. **Table 4** lists the feature families, which correspond to the method groups in Section 4.

Table 4. Feature families for continuous identity evaluation.

Feature family	Examples	Primary question	Failure mode to control
Authentication context	Hour, source, geovelocity, device, protocol, MFA result, token age	Is session establishment unusual for this identity or peer group?	Travel, remote work, and infrastructure changes create benign novelty
Resource behaviour	First-seen host, application, file, repository, data sensitivity	Is the identity accessing atypical or high-value resources?	Legitimate projects change resource needs
Privilege behaviour	Role assumption, group change, delegated scope, administrative tools	Is authority expanding or being exercised unusually?	Administrators legitimately have broad variance
Temporal sequence	Event order, inter-event time, burstiness, session duration	Does the trajectory resemble adversary progression?	Sparse logging breaks sequences
Graph structure	New edges, path rarity, fan-out, centrality, cross-zone movement	Is the identity creating unusual relationships?	Topology evolves; future-edge leakage inflates results
Process and network context	Process tree, remote-service use, DNS, egress, bytes	Do endpoint and network actions corroborate identity risk?	NAT, encryption, and shared systems weaken attribution

7.2. Risk Fusion

For identity i at time t, the framework estimates continuous session risk as

$$R(i, t) = \sigma[w_1 A(i, t) + w_2 P(i, t) + w_3 G(i, t) + w_4 S(i, t) + w_5 C(i, t) - w_6 Q(i, t)] \quad (1)$$

where A is the authentication-context anomaly, P the personal and peer deviation, G the graph anomaly, S the temporal-sequence anomaly, C the ATT&CK-consistency and cross-source corroboration term, Q a data-quality or uncertainty penalty, w the learned or policy-calibrated weights, and σ a monotone map to [0, 1]. Two terms in Equation (1) are motivated by the review. The term C rewards evidence chains that match a plausible ATT&CK path rather than isolated anomalies, addressing the false-positive problem identified in peer-group work [33]. The term Q lowers confidence when endpoint coverage is missing or identity resolution is ambiguous, so that a high raw score on poorly observed telemetry does not trigger an aggressive response. Weights and σ should be calibrated on a validation period so that R approximates the probability that a session is adversary-controlled at the operating alert budget.

7.3. ATT&CK-Aware Explanation and Response Policy

Each alert should expose the evidence chain, candidate ATT&CK techniques, confidence, affected resources, and the smallest defensible response. For example: "U1847 authenticated from a previously unseen source, queried directory groups, created three new remote-service edges, and initiated an abnormal outbound transfer. Candidate path: T1078 → T1087/T1069 → T1021 → T1567; confidence 0.82." The word *candidate* avoids claiming that telemetry proves intent. Attack-chain reconstruction on LANL data [46] shows that such chains can be assembled from heterogeneous logs.

Risk should drive graduated controls: low-confidence novelty may trigger observation; moderate risk may require step-up authentication or reduced token lifetime; high-confidence multi-source sequences may restrict privileges, revoke sessions, isolate devices, or open an investigation. Automated containment should require stronger corroboration than analyst notification.

8. Reproducible Empirical Validation Protocol

The LANL Comprehensive, Multi-Source Cyber-Security Events dataset is the recommended primary benchmark: it contains 1,648,275,307 authentication, process, DNS, network-flow, and labelled red-team events for 12,425 users and 17,684 computers over 58 days [15]. The nine-month LANL authentication dataset (708,304,516 user-computer associations) [65], the unified host and network dataset [16], and the CERT releases [21] play supporting roles (**Table 5**).

The protocol has seven steps. (1) Build a canonical event table and resolve user, source, destination, process, and session identifiers without future information. Because LANL records carry no session identifier, a session is the ordered set of events for one user and source computer from a successful logon to the matching

logoff or, failing that, an inactivity gap fixed on the training window; process, DNS, and flow events are attached by user or computer and timestamp. A session is positive when it contains a red-team event matched on time, user, source, and destination; sessions adjacent to a labelled event are also reported separately, identity and campaign labels are aggregated from session labels, and sensitivity to the gap is reported. (2) Use chronological training, validation, and test windows, never random event splits, because adjacent events from one session leak across partitions. (3) Construct baselines: a frequency or rule model, an isolation-forest or robust statistical detector, a peer-group tensor model [17], a sequence model [37], a temporal graph model [18] [27], and the full ATT&CK-aware fusion model. (4) Measure event-, session-, identity-, and campaign-level performance and time to first detection relative to the first labelled red-team event. (5) Tune thresholds only on the validation period under explicit alert budgets, such as alerts per 10,000 users per day. (6) Run ablations removing authentication context, process events, network events, graph features, sequence features, ATT&CK corroboration, and the uncertainty penalty. (7) Test robustness to missing telemetry, delayed events, new employees, role changes, service accounts, and drift [40].

Table 5. Candidate datasets and their roles.

Dataset	Strength	Limitation	Recommended use
LANL multi-source (58 days)	Cross-source telemetry plus red-team authentication labels	Short period; anonymised context; labels cover selected actions	Primary benchmark and ablation study
LANL authentication associations (9 months)	Large scale and long temporal horizon	Successful authentication triples only; no attack ground truth	Unsupervised scalability and drift analysis
LANL unified host and network (90 days)	Joined host and network views	No attack labels	Entity resolution and corroboration development
CERT insider threat (r4.2, r5.2, r6.2)	Multi-modal user actions and scenario labels	Synthetic behaviour; insider focus	Transfer test for malicious authorised users
DARPA Transparent Computing and OpTC	Provenance-level host telemetry with red-team ground truth	Small user population; research instrumentation	Provenance corroboration and path reconstruction
Organisation-specific audit logs	Real policies, resources, and response costs	Privacy, access, reproducibility, incomplete labels	External validation after governance approval

Accuracy and AUROC are incomplete evidence under extreme imbalance. Primary metrics should include the area under the precision-recall curve, precision at a fixed alert budget, recall at fixed false-positive rates, alerts per analyst-day,

mean time to detection, and campaign coverage. Calibration (reliability diagrams, Brier score) and day- or campaign-level bootstrap confidence intervals should also be reported. Table 6 specifies the evidence required for each major claim.

Table 6. Evidence required for major empirical claims.

Claim	Required comparison	Minimum evidence
Cross-source context improves detection	Full model versus authentication-only model	Chronological test; same alert budget; bootstrap confidence interval
ATT&CK explanations are faithful	Mapped explanation versus event evidence	Analyst review; technique-level precision and recall; disagreement report
Model works at operational scale	Throughput and memory by events and users	Wall-clock benchmark; reproducible hardware configuration
Risk scores support response	Calibrated versus uncalibrated scores	Brier score, reliability plot, utility or cost analysis
Performance generalises	LANL versus CERT, OpTC, or an external site	Frozen model or documented adaptation; shift analysis

9. Critical Challenges and Open Research Questions

Ground truth and base rates. Red-team labels identify known attack actions but do not establish that every unlabelled event is benign, so a detector that flags preparatory or consequential behaviour may be charged a false positive. Evaluation should use temporal tolerance windows, campaign-level review, and positive-unlabelled assumptions; synthetic injection [36] extends coverage but cannot establish performance against authentic adversaries. At enterprise scale even a low false-positive rate can produce an unmanageable queue [47]: a dual-granularity model that produced 26 daily anomalies against 64 from a baseline on 15,000 users [44] shows why counts per analyst-day should be reported alongside rates.

Drift and identity diversity. Static baselines become stale as behaviour changes, while rapid adaptation lets slow attacks normalise. Bounded adaptation, change-point detection [23], drift-triggered retraining [40], and policy-aware peer groups need study. Non-human identities do not follow human circadian or geolocation patterns, yet only one included study concentrated on them.

Privacy and mapping uncertainty. Behavioural monitoring should be purpose-limited and minimised; differentially private UEBA [66] and federated insider detection [31] show that privacy and detection can coexist. Because ordinary administration produces the same observables as ATT&CK techniques, mappings should combine rules, probabilistic evidence, and analyst feedback, and report candidate techniques with confidence.

10. Implications for Critical Infrastructure

In critical infrastructure, IT and OT identities traverse remote-access gateways, jump hosts, engineering workstations, and vendor accounts, while availability and

safety constraints limit automated containment. A defensible implementation needs sector-specific telemetry mappings, high-confidence response tiers, break-glass procedures, and human approval for actions that could interrupt essential services; ATT&CK for ICS should be used when a sequence reaches control-system objectives [67]. Complementary work addresses adjacent layers: generative modelling for vulnerability discovery and infrastructure hardening [68], blockchain-based federated identity management for cross-sector trust [69], anomaly detection over infrastructure sensor networks [70], and decentralised policy enforcement and anomaly response that removes the central policy engine as a single point of failure [71]. The proposed framework sits between those layers: it consumes the identity evidence that federated identity management produces and supplies the session-level risk that sensor- and vulnerability-focused defences do not observe.

11. Discussion

Post-authentication detection is a multi-source evidence-fusion problem. The strongest studies, which model authentication topology or sequence on LANL data with chronological evaluation [17] [18] [26] [27] [37], support the hypothesis that conditioning risk on the evolving session reduces the false alarms of isolated login novelty, while also showing how low precision remains at realistic base rates. ATT&CK supplies adversary semantics and a common vocabulary [20] [46] but not organisational baselines, event schemas, probabilistic estimation, or response-cost models, so the proposed architecture uses it as a corroboration and explanation layer above statistical detection. Its central hypothesis is testable: at the same alert budget, combining post-login sequence and graph evidence with authentication context will detect more compromised-identity campaigns, earlier, than authentication-only models. Secondary hypotheses concern explanation fidelity, calibration, and robustness to missing telemetry and to adversarial evasion, which sharply degrades graph-based fraud detectors without adversarial training [72].

12. Limitations

Records came from a single index using phrase-based queries, so studies using other terminology may have been missed, although citation searching mitigated this risk; the computer-science filter further limits coverage of cloud-governance, organisational, and critical-infrastructure research. Screening was completed by a single reviewer with rule-assisted classification, which introduces selection risk despite the preserved decision record. Extraction for most studies was at abstract level, so the percentages in Section 5.1 are lower bounds, and eligibility for the studies coded from title and abstract alone was not verified in full text. The review did not itself compare cross-source with authentication-only detection, so that comparison remains a hypothesis for the protocol in Section 8. The risk formulation is conceptual, and public enterprise datasets are anonymised, dated, and incomplete.

13. Conclusion

Identity assurance must continue throughout the authenticated session. This review of 135 studies shows a rapidly expanding literature whose strongest evidence comes from graph, sequence, and tensor models evaluated chronologically on the LANL datasets, and whose principal weaknesses are the scarcity of chronological validation, realistic class-balance reporting, operational alert metrics, and adversary-aligned explanation. Empirical implementation on the LANL multi-source data, followed by transfer testing and sector-specific validation, is required to establish effectiveness and deployment suitability.

Author Contributions

Samuel Dushimimana Shyaka and Collin Kabwama: Conceptualization, Methodology, Investigation, Data curation, Formal analysis, Visualization, Writing (original draft), Writing (review and editing); Dushime Benjamin: Data curation, Formal analysis, Visualization, Writing (review and editing).

Data Availability

No new primary data were created. The search log, screening decisions with reason codes, extraction table, and appraisal scores are provided as supplementary tables S1 to S3, which are not printed in the article but are supplied as a separate supplementary workbook available at https://zenodo.org/records/23284504. The proposed validation uses the publicly documented datasets cited in Section 8.

Conflicts of Interest

The authors declare no conflicts of interest.

References

[1] Verizon (2025) 2025 Data Breach Investigations Report. Verizon Business. https://www.verizon.com/business/resources/reports/dbir/

[2] Microsoft (2025) Microsoft Digital Defense Report 2025. Microsoft Corporation. https://www.microsoft.com/en-us/corporate-responsibility/cybersecurity/microsoft-digital-defense-report-2025/

[3] MITRE (2026) Valid Accounts, Technique T1078. MITRE ATT&CK for Enterprise. https://attack.mitre.org/techniques/T1078/

[4] Shah, N., Ho, G., Schweighauser, M., Ibrahim, M., Cidon, A. and Wagner, D. (2020) A Large-Scale Analysis of Attacker Activity in Compromised Enterprise Accounts. In: Wang, G., Ciptadi, A. and Ahmadzadeh, A. Eds., *Communications in Computer and Information Science*, Springer International Publishing, 3-27. https://doi.org/10.1007/978-3-030-59621-7_1

[5] Rose, S., Borchert, O., Mitchell, S. and Connelly, S. (2020) Zero Trust Architecture. National Institute of Standards and Technology.

[6] Cybersecurity and Infrastructure Security Agency (2023) Zero Trust Maturity Model, Version 2.0. CISA. https://www.cisa.gov/zero-trust-maturity-model

[7] Gheyas, I.A. and Abdallah, A.E. (2016) Detection and Prediction of Insider Threats

to Cyber Security: A Systematic Literature Review and Meta-Analysis. *Big Data Analytics*, **1**, Article No. 6. https://doi.org/10.1186/s41044-016-0006-0

[8] Gong, Y., Cui, S., Liu, S., Jiang, B., Dong, C. and Lu, Z. (2024) Graph-Based Insider Threat Detection: A Survey. *Computer Networks*, **254**, Article 110757. https://doi.org/10.1016/j.comnet.2024.110757

[9] Landauer, M., Onder, S., Skopik, F. and Wurzenberger, M. (2023) Deep Learning for Anomaly Detection in Log Data: A Survey. *Machine Learning with Applications*, **12**, Article 100470. https://doi.org/10.1016/j.mlwa.2023.100470

[10] Al-Sada, B., Sadighian, A. and Oligeri, G. (2025) MITRE ATT&CK: State of the Art and Way Forward. *ACM Computing Surveys*, **57**, 1-37. https://doi.org/10.1145/3687300

[11] Adewole, K.S., Anuar, N.B., Kamsin, A., Varathan, K.D. and Razak, S.A. (2017) Malicious Accounts: Dark of the Social Networks. *Journal of Network and Computer Applications*, **79**, 41-67. https://doi.org/10.1016/j.jnca.2016.11.030

[12] P. Velayudhan, S. and Somasundaram, M.S.B. (2019) Compromised Account Detection in Online Social Networks: A Survey. *Concurrency and Computation*: *Practice and Experience*, **31**, e5346. https://doi.org/10.1002/cpe.5346

[13] Page, M.J., McKenzie, J.E., Bossuyt, P.M., Boutron, I., Hoffmann, T.C., Mulrow, C.D., *et al.* (2021) The PRISMA 2020 Statement: An Updated Guideline for Reporting Systematic Reviews. *BMJ*, **372**, n71. https://doi.org/10.1136/bmj.n71

[14] Kitchenham, B., Pearl Brereton, O., Budgen, D., Turner, M., Bailey, J. and Linkman, S. (2009) Systematic Literature Reviews in Software Engineering—A Systematic Literature Review. *Information and Software Technology*, **51**, 7-15. https://doi.org/10.1016/j.infsof.2008.09.009

[15] Kent, A.D. (2016) Cyber Security Data Sources for Dynamic Network Research. In: Hankin, C., Ed., *Security Science and Technology*, World Scientific (Europe), 37-65. https://doi.org/10.1142/9781786340757_0002

[16] Turcotte, M.J.M., Kent, A.D. and Hash, C. (2018) Unified Host and Network Data Set. In: Hankin, C., Ed., *Security Science and Technology*, World Scientific (Europe), 1-22. https://doi.org/10.1142/9781786345646_001

[17] Eren, M.E., Moore, J.S., Skau, E., Moore, E., Bhattarai, M., Chennupati, G., *et al.* (2023) General-Purpose Unsupervised Cyber Anomaly Detection via Non-Negative Tensor Factorization. *Digital Threats*: *Research and Practice*, **4**, 1-28. https://doi.org/10.1145/3519602

[18] King, I.J. and Huang, H.H. (2023) Euler: Detecting Network Lateral Movement via Scalable Temporal Link Prediction. *ACM Transactions on Privacy and Security*, **26**, 1-36. https://doi.org/10.1145/3588771

[19] Liu, F., Wen, Y., Zhang, D., Jiang, X., Xing, X. and Meng, D. (2019) Log2vec: A Heterogeneous Graph Em-Bedding Based Approach for Detecting Cyber Threats within Enterprise. *Proceedings of the* 2019 *ACM SIGSAC Conference on Computer and Communications Security*, London, 11-15 November 2019, 1777-1794. https://doi.org/10.1145/3319535.3363224

[20] Serzantov, V., Kupris, E. and Schreck, T. (2025) Identity Threats and Where to Find Them: Mapping ITDR and MITRE ATT&CK. 2025 *IEEE European Symposium on Security and Privacy Workshops* (*EuroS&PW*), Venice, 30 June 2025-4 July 2025, 485-495. https://doi.org/10.1109/eurospw67616.2025.00062

[21] Glasser, J. and Lindauer, B. (2013) Bridging the Gap: A Pragmatic Approach to Generating Insider Threat Data. 2013 *IEEE Security and Privacy Workshops*, San Francisco, 23-24 May 2013, 98-104. https://doi.org/10.1109/spw.2013.37

[22] Egele, M., Stringhini, G., Kruegel, C. and Vigna, G. (2017) Towards Detecting Compromised Accounts on Social Networks. *IEEE Transactions on Dependable and Secure Computing*, **14**, 447-460. https://doi.org/10.1109/tdsc.2015.2479616

[23] Wu, T. and Shan, M. (2025) Exploring Compromised Accounts in Social Media by Change Point Detection. In: Chu, WT., Shen, CY. and Shuai, HH., Eds., *Communications in Computer and Information Science*, Springer, 260-274. https://doi.org/10.1007/978-981-96-4596-1_20

[24] Siadati, H. and Memon, N. (2017) Detecting Structurally Anomalous Logins within Enterprise Networks. *Proceedings of the* 2017 *ACM SIGSAC Conference on Computer and Communications Security*, Dallas, 30 October-3 November 2017, 1273-1284. https://doi.org/10.1145/3133956.3134003

[25] Kaiafas, G., Varisteas, G., Lagraa, S., State, R., Nguyen, C.D., Ries, T., *et al.* (2018) Detecting Malicious Authentication Events Trustfully. *NOMS* 2018-2018 *IEEE/IFIP Network Operations and Management Symposium*, Taipei, China, 23-27 April 2018, 1-6. https://doi.org/10.1109/noms.2018.8406295

[26] Tang, M. and Ge, B. (2025) Social-Hunter: A Social Heuristics-Based Approach to Early Unveiling Unknown Malicious Logins Using Valid Accounts. *Computers & Security*, **150**, Article 104269. https://doi.org/10.1016/j.cose.2024.104269

[27] Bowman, B., Laprade, C., Ji, Y. and Huang, H.H. (2020) Detecting Lateral Movement in Enterprise Computer Networks with Unsupervised Graph AI. *Proceedings of the* 23*rd International Symposium on Research in Attacks, Intrusions and Defenses* (*RAID* 2020), San Sebastián, Spain, 14-16 October 2020, 257-268.

[28] Bian, H., Bai, T., Salahuddin, M.A., Limam, N., Daya, A.A. and Boutaba, R. (2021) Uncovering Lateral Movement Using Authentication Logs. *IEEE Transactions on Network and Service Management*, **18**, 1049-1063. https://doi.org/10.1109/tnsm.2021.3054356

[29] Sharma, B., Pokharel, P. and Joshi, B. (2020) User Behavior Analytics for Anomaly Detection Using LSTM Autoencoder-Insider Threat Detection. *Proceedings of the* 11*th International Conference on Advances in Information Technology*, Bangkok, 1-3 July 2020, 1-9 https://doi.org/10.1145/3406601.3406610

[30] Paduvilan, A.K., Babu, S.A., Sekar, J., Elengovan, A.K., Seshagiri, N., Jangiti, K., *et al.* (2026) Attention-Driven Fusion of Sequential and Graph Neural Models for Insider Threat Detection in UEBA. *Scientific Reports*, **16**, Article No. 21117. https://doi.org/10.1038/s41598-026-41242-7

[31] Godson Samuel, J.J., Balasubramaniya, V., Kumar, G.G., Geetha, R. and Shobana, R. (2026) Insider Threat Detection System with Federated Learning and Differential Privacy. 2026 *International Conference on Intelligent Computing, Networks, and Security* (*IC-ICNS*), Bhubaneswar, 26-28 March 2026, 1-6. https://doi.org/10.1109/ic-icns68863.2026.11537834

[32] Cheng, K., Li, D., Tang, W., Shu, Y. and Qiang, W. (2026) Insider Threat Detection and Evidence-Grounded Triage Using a Heterogeneous Dual-Branch Network and a Constrained LLM Agent. *Applied Sciences*, **16**, Article 7054. https://doi.org/10.3390/app16147054

[33] Tang, B., Hu, Q. and Lin, D. (2017) Reducing False Positives of User-to-Entity First-Access Alerts for User Behavior Analytics. 2017 *IEEE International Conference on Data Mining Workshops* (*ICDMW*), New Orleans, 18-21 November 2017, 804-811. https://doi.org/10.1109/icdmw.2017.111

[34] Ibadildin, N., Batkuldin, A., Bekmurat, N. and Khanezar, S. (2026) Zero-Trust Architecture with Continuous AI-Based Risk Scoring for Dynamic Access Control. 2026

International Conference on Cybersecurity, Digital Forensics, and AI Applications (*ICCSDFAI*), Istanbul, 25-27 June 2026, 1183-1188. https://doi.org/10.1109/iccsdfai70505.2026.11647932

[35] Goodman, E., Ingram, J., Martin, S. and Grunwald, D. (2015) Using Bipartite Anomaly Features for Cyber Security Applications. 2015 *IEEE* 14*th International Conference on Machine Learning and Applications* (*ICMLA*), Miami, 9-11 December 2015, 301-306. https://doi.org/10.1109/icmla.2015.69

[36] Muratović, L. and Isaković, B. (2026) A Comparative Study of Rule-Based and Machine Learning-Based Methods for Authentication Anomaly Detection. 2026 30*th International Conference on Information Technology* (*IT*), Zabljak, 24-28 February 2026, 1-4. https://doi.org/10.1109/it67293.2026.11435797

[37] Gökstorp, S., Nyberg, J., Kim, Y., Johnson, P. and Dán, G. (2024) Anomaly Detection in Security Logs Using Sequence Modeling. *NOMS* 2024-2024 *IEEE Network Operations and Management Symposium*, Seoul, 6-10 May 2024, 1-9. https://doi.org/10.1109/noms59830.2024.10575561

[38] Kabwama, C.A., Businge, P., Malingu, C.J., *et al.* (2025) Graph Attention Networks for Credit Card Fraud Detection: A Relational Learning Approach. *World Journal of Advanced Research and Reviews*, **26**, 2580-2585. https://doi.org/10.30574/wjarr.2025.26.3.2400

[39] Businge, P., Agaba, I.A., Atuhaire, J.I., Nayebale, F.I., Ariho, J.G., Mugalu, B., *et al.* (2025) Adaptive Financial Fraud Detection Using Graph Neural Networks and Reinforcement Learning. *World Journal of Advanced Research and Reviews*, **28**, 842-847. https://doi.org/10.30574/wjarr.2025.28.2.3761

[40] Abdulhussein, I.M. (2026) Graph Neural Network for Detecting Lateral Movement in Enterprise Logs with Concept Drift. *Scientia. Technology, Science and Society*, **3**, 75-88. https://doi.org/10.59324/stss.2026.3(3).06

[41] Nina, P.P., Paang, M. and Reiglenc, J. (2026) Zero-Trust Graph Neural Networks: Dynamic Lateral Movement Detection and Attack Path Reconstruction in Hybrid Cloud Infrastructure. *International Journal of Applied Engineering and Intelligent Computing*, **1**, 10-13. https://doi.org/10.66917/ijaeic.a000014

[42] Goud, K.N., Jain, K. and Krishnan, P. (2026) Identity-Centric Detection of Advanced Persistent Threats via Privilege Transition Graphs. 2026 6*th International Conference on Inventive Computation and Information Technologies* (*ICICIT*), Salem, 6-8 July 2026, 1702-1707. https://doi.org/10.1109/icicit69063.2026.11633950

[43] Aljumaily, M.S. and Abd, H. (2025) Enhancing User and Entity Behavior Analytics in SIEM Systems Using AI-Powered Anomaly Detection: A Data-Driven Simulation Approach. *International Journal of Mechatronics, Robotics, and Artificial Intelligence*, **1**, 82-94. https://doi.org/10.33971/ijmrai.1.2.11

[44] Jordan, A. and Chen, Y. (2025) User and Entity Behavior Analytics (UEBA) Enhanced Security Anomaly Detection in Enterprise DevSecOps Platforms. 2025 *IEEE Secure Development Conference* (*SecDev*), Indianapolis, 14-16 October 2025, 94-104. https://doi.org/10.1109/secdev66745.2025.00021

[45] Malingu, C.J., Kabwama, C.A., Businge, P., *et al.* (2025) Application of LLMS to Fraud Detection. *World Journal of Advanced Research and Reviews*, **26**, 178-183. https://doi.org/10.30574/wjarr.2025.26.2.1586

[46] Zhang, X., Wang, Y., Zhang, C. and Chi, J. (2025) Research on Attack Chain Reconstruction Technology Based on Graph Analysis, Machine Learning, and Threat Intelligence: Application of Graph Analysis and Machine Learning in Attack Chain Reconstruction. *Proceedings of the* 2025 6*th International Conference on Blockchain*

Technology, Huhehaote, 15-17 August 2025, 30-38. https://doi.org/10.1145/3778335.3778346

[47] Axelsson, S. (2000) The Base-Rate Fallacy and the Difficulty of Intrusion Detection. *ACM Transactions on Information and System Security*, **3**, 186-205. https://doi.org/10.1145/357830.357849

[48] Artioli, P., Maci, A. and Magrì, A. (2024) A Comprehensive Investigation of Clustering Algorithms for User and Entity Behavior Analytics. *Frontiers in Big Data*, **7**, Article ID: 1375818. https://doi.org/10.3389/fdata.2024.1375818

[49] Siadati, H., Saket, B. and Memon, N. (2016) Detecting Malicious Logins in Enterprise Networks Using Visualization. 2016 *IEEE Symposium on Visualization for Cyber Security* (*VizSec*), Baltimore, 24 October 2016, 1-8. https://doi.org/10.1109/vizsec.2016.7739582

[50] Kodumuri, S. and Dadi, R. (2025) Deep Learning-Enhanced Multi-Modal Fusion for Anomaly-Based Insider Threat Detection in Enterprise Environments. 2025 *International Conference on NexGen Networks and Cybernetics* (*IC2NC*), Erode, 1-3 December 2025, 188-194. https://doi.org/10.1109/ic2nc67409.2025.11376104

[51] Mohideen, A.B., Krishna Sai, P., Basha, Y.S., Sohel, S. and Reddy, Y.V. (2025) Enhancing Endpoint Security through AI Powered Anomaly Detection in Endpoint Logs Using Local Outlier Factor (LOF). 2025 6*th International Conference on Intelligent Communication Technologies and Virtual Mobile Networks* (*ICICV*), Tirunelveli, 17-19 June 2025, 1403-1411. https://doi.org/10.1109/icicv64824.2025.11085490

[52] Amrouche, F., Lagraa, S., Kaiafas, G. and State, R. (2019) Graph-Based Malicious Login Events Investigation. 2019 *IFIP/IEEE Symposium on Integrated Network and Service Management* (*IM*), Arlington, 8-12 April 2019, 63-66. https://ieeexplore.ieee.org/document/8717788

[53] Sunny, M.A.S., Ahmed, I., Arman, H.U. and Rahman, M. (2025) Integrating UEBA-Driven Risk Scoring for Adaptive Attribute Based Decryption. 2025 *IEEE 4th International Conference on Robotics, Automation, Artificial-Intelligence and Internet-of-Things* (*RAAICON*), Dhaka, 27-28 November 2025, 680-685. https://doi.org/10.1109/raaicon69033.2025.11502528

[54] Eren, M.E., Moore, J.S. and Alexandro, B.S. (2020) Multi-Dimensional Anomalous Entity Detection via Poisson Tensor Factorization. 2020 *IEEE International Conference on Intelligence and Security Informatics* (*ISI*), Arlington, 9-10 November 2020, 1-6. https://doi.org/10.1109/isi49825.2020.9280524

[55] Kaiafas, G., Hammerschmidt, C., Lagraa, S. and State, R. (2020) Auto Semi-Supervised Outlier Detection for Malicious Authentication Events. In: Cellier, P. and Driessens, K., Eds., *Communications in Computer and Information Science*, Springer International Publishing, 176-190. https://doi.org/10.1007/978-3-030-43887-6_14

[56] Thakare, A.V., Nawale, I.G., Deole, B.P., Bobade, S.R. and Dhopte, S.M. (2026) AI-Based Insider Threat Detection in Cloud. *International Journal of Scientific Research in Engineering and Management*, **10**, 1-9. https://doi.org/10.55041/ijsrem56881

[57] Wei, R., Cai, L., Yu, A. and Meng, D. (2020) AGE: Authentication Graph Embedding for Detecting Anomalous Login Activities. In: Zhou, J., Luo, X., Shen, Q. and Xu, Z., Eds., *Lecture Notes in Computer Science*, Springer International Publishing, 341-356. https://doi.org/10.1007/978-3-030-41579-2_20

[58] Maltsev, V. and Budovskaya, G. (2026) Comparative Analysis of Digital Footprint Sources for UEBA-Class Systems with Quantitative Assessment of Informativeness and Computational Cost. *Proceedings of the International Scientific and Practical*

Conference, Voronezh, 30 March 2026, 963-971.
https://doi.org/10.58168/mist2026_963-971

[59] Zhao, S., Wei, R., Cai, L., Yu, A. and Meng, D. (2020) CTLMD: Continuous-Temporal Lateral Movement Detection Using Graph Embedding. In: Zhou, J., Luo, X., Shen, Q. and Xu, Z., Eds., *Lecture Notes in Computer Science*, Springer International Publishing, 181-196. https://doi.org/10.1007/978-3-030-41579-2_11

[60] Mozumder, M.A.S., Mia, M.M., Nipa, R.A., Anik, A., Bhuiyan, K.R., Akter, S., *et al.* (2026) A Hybrid Graph Neural Network and Large Language Model Framework for Insider Threat Detection via Behavioral Graph and Semantic Profiling. *International Journal of Computer Science & Information System*, **11**, 22-32.
https://doi.org/10.55640/ijcsis/volume11issue05-03

[61] Wu, Y., Liu, F. and Wen, Y. (2021) Malicious Login Detection Using Long Short-Term Memory with an Attention Mechanism. In: Peterson, G. and Shenoi, S., Eds., *IFIP Advances in Information and Communication Technology*, Springer International Publishing, 157-173. https://doi.org/10.1007/978-3-030-88381-2_8

[62] Athapaththu, A.M.M.I.P., Rathnayaka, R.M.C.A., Wanasinghe, W.M.K.R., *et al.* (2026) A Hybrid Multi-Modal AI Framework for Insider Threat Detection in Hybrid Work Environments. 2026 8*th International Congress on Human-Computer Interaction, Optimization and Robotic Applications* (*ICHORA*), Ankara, 21-23 May 2026, 1-10. https://doi.org/10.1109/ichora69329.2026.11537058

[63] Thammaiah, N., Girish, N., Meghana, N., G., S. and Kenny, S. (2026) An Interpretable Insider Threat Detection Framework: Correlating Digital Communications and Behavioral Metadata. *Proceedings of the* 18*th International Conference on Agents and Artificial Intelligence*, **1**, 864-871. https://doi.org/10.5220/0014685300004052

[64] Trinita, V.S. and Kurniawan, A. (2026) Fortifying the SOC: A Validated Lightgbm Engine for Behavioral Insider Threat Detection. 2026 4*th International Conference on Intelligent Data Communication Technologies and Internet of Things* (*IDCIoT*), Ottapalam, 4-6 February 2026, 245-251.
https://doi.org/10.1109/idciot67589.2026.11455760

[65] Kent, A.D. (2014) User-Computer Authentication Associations in Time. Los Alamos National Laboratory.

[66] Rashid, F. and Miri, A. (2021) User and Event Behavior Analytics on Differentially Private Data for Anomaly Detection. 2021 7*th IEEE International Conference on Big Data Security on Cloud* (*BigDataSecurity*), *IEEE International Conference on High Performance and Smart Computing*, (*HPSC*) *and IEEE International Conference on Intelligent Data and Security* (*IDS*), New York, 15-17 May 2021, 81-86.
https://doi.org/10.1109/bigdatasecurityhpscids52275.2021.00025

[67] MITRE (2026) ATT&CK for Industrial Control Systems.
https://attack.mitre.org/matrices/ics/

[68] Kabwama, C.A., Ayozie, O.M., Zeyeum, J.N., Awe, A.O., Ikwuogu, O.F. and Titilayo, D.H. (2025) Generative Adversarial Network (GAN) Modeling for CNI Vulnerability Discovery: An Immutable Knowledge Base for Automated Infrastructure Hardening via Distributed Ledger Technology. *Global Journal of Engineering and Technology Advances*, **25**, 292-307. https://doi.org/10.30574/gjeta.2025.25.1.0309

[69] Pinyi, E.O., Kabwama, C.A., Zeyeum, J.N., Oluwabukola, H.P. and Ikwuogu, O.F. (2025) Federated IAM for Critical Systems: A Blockchain-Based Decentralized Identity Framework Enhanced by Federated Learning for Cross-Sector CNI Trust. *International Journal of Computer Applications Technology and Research*, **14**, 28-45.
https://doi.org/10.7753/ijcatr1409.1005

[70] Olugbade, A., Adediji, E., Eni, F., Titilayo, D.H., Kabwama, C.A. and Agbesi, J.S. (2025) Hybrid Quantum Machine Learning for Anomaly Detection in Critical Infrastructure Sensor Networks. *Global Journal of Engineering and Technology Advances*, **23**, 333-361. https://doi.org/10.30574/gjeta.2025.23.3.0198

[71] Kabwama, C.A., Pinyi, E.O., Adediji, E., Zeyeum, J.N. and Ikwuogu, O.F. (2025) Decentralized AI-Driven Zero-Trust Architecture: Leveraging Blockchain for Immutable Policy Enforcement and Autonomous Anomaly Response in Critical Infrastructure Systems. *World Journal of Advanced Engineering Technology and Sciences*, **17**, 559-573. https://doi.org/10.30574/wjaets.2025.17.1.1427

[72] Malingu, C.J., Titilayo, D.H., Eni, F., Kabwama, C.A. and Anyah, V. (2024) Adversarial Machine Learning in Finance: Developing Resilient AI Models to Counter Fraudster Evasion Attacks on US Bank Security Systems. *International Journal of Computer Applications Technology and Research*, **13**, 111-137.