From Alert Flood to Actionable Intelligence: Reconstructing Attack Chains via TTP Sequences Using LLMs and RAG

Ruijie Qi[a,b], Wenxin Le[a,b], Ruiqi Wang[a⋆], Yingxiao Xiang[a], Yepeng Yao[a,b], and Zhengwei Jiang[a,b]

[a]Institute of Information Engineering, Chinese Academy of Sciences, Beijing, China
[b]School of Cyber Security, University of Chinese Academy of Sciences, Beijing, China
{qiruijie, lewenxin, wangruiqi, xiangyingxiao, yaoyepeng, jiangzhengwei}@iie.ac.cn

Abstract. Modern intrusion detection systems (IDSs) in large-scale enterprise networks generate massive volumes of heterogeneous alerts, many of which are weakly correlated interaction traces that obscure the semantic progression of multi-stage attacks. While existing alert filtering techniques reduce analyst workload, they often fail to preserve behavioral continuity necessary for reconstructing coherent attack campaigns.
In this paper, we propose a TTP-Guided Event Reconstruction Framework that integrates Cyber Threat Intelligence (CTI) with alert-side interaction evidence to reconstruct semantically coherent attack events from noisy intrusion alerts. Structured ATT&CK-aligned TTP sequences are extracted from CTI reports to form a knowledge space of canonical attack progressions. Candidate TTP paths derived from alerts are embedded into a shared semantic space, filtered using a one-class representation model, and aggregated into candidate events. A retrieval-augmented reconstruction module further refines attacker–victim structures by aligning alert-derived behaviors with CTI-derived progression templates. Experiments on three publicly available cyber range datasets (CPTC-2017, CPTC-2018, and CCDC-2018) show that the proposed framework reduces structurally inconsistent interaction paths while improving attacker–victim coherence at the event level. Reconstructed events exhibit improved semantic alignment with CTI-derived attack progression patterns compared to their originating alert paths.

Keywords: Cyber Threat Intelligence · MITRE ATT&CK · Alert Noise Reduction · Large Language Model–Driven Analysis.

1 Introduction

The cybersecurity capability of Security Operations Centers (SOCs) largely depends on network situational awareness derived from alert data collected from heterogeneous security devices [1]. Leveraging techniques like statistical analysis,

⋆ Corresponding Author: Ruiqi Wang

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

correlation, and AI-based analysis on these data enables threat analysis [2], system risk assessment [3], and incident management [4]. However, real-world alert streams typically contain substantial noise, which obscures genuine security indicators and hinders the accurate detection of high-risk events [5]. Effective noise reduction is therefore essential for maintaining reliable and actionable security operations. Traditional rule-based filtering approaches, however, are difficult to maintain and adapt due to the continual evolution of attack techniques, making it imperative to explore LLM-driven adaptive noise reduction frameworks.

Although prior research [6, 7] has made significant progress in understanding and mitigating noise alerts, these studies still fail to address the underlying causes. Existing methods insufficiently account for factors such as pervasive middleware deployment, heterogeneous detection capabilities across devices, and increasingly sophisticated attacker behaviors. As a result, false positives and redundant alerts remain pervasive.

In summary, our contributions are as follows:

- We propose a TTP-Guided Event Reconstruction Framework that integrates CTI-derived attack knowledge with alert-side interaction evidence for coherent attack-event reconstruction.
- Within this framework, we design a Semantic Knowledge Layer that represents attack behaviors as structured TTP sequences extracted from open-source threat intelligence reports, providing a progression-aware semantic prior for interpreting alert-derived evidence.
- We further develop an Attack Modeling Layer and an Event Reconstruction Layer that respectively perform candidate TTP path generation with structural filtering and retrieval-augmented event refinement, improving event coherence and semantic alignment on three public cyber range datasets.

2 Related work

The efficient analysis of security alerts is fundamental to modern cybersecurity operations [8]. Research in this domain has branched into two interconnected directions: techniques for prioritizing alerts to reduce analyst burden, and methods for extracting structured knowledge from raw alert streams to provide context.

2.1 Alert Prioritisation

Alert prioritisation aims to separate benign events from genuine threats and rank security alerts by their severity and relevance, directly addressing the issue of alert fatigue. Early approaches often relied on rule-based systems and expert knowledge to assign severity scores, but these methods struggle to adapt to evolving attack patterns.

To address this problem, statistical and machine learning models have been widely adopted. For instance, probabilistic models such as the Hidden Markov Model (HMM) have been used to model attack stages and infer the criticality of subsequent alerts [9]. Similarly, sequence-to-sequence architectures based

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

on LSTM networks capture long-range temporal dependencies to predict the progression of an attack, thereby enabling proactive prioritisation of related alerts [10].

More recent research emphasizes operational augmentation for analysts. Point-cloud clustering combined with machine learning classifiers supports automated separation of genuine threats from benign activities, streamlining Security Orchestration, Automation and Response workflows [6]. Addressing the core problem of alert fatigue directly, reinforcement-learning–driven frameworks such as AlertPro dynamically learn optimal prioritisation policies, adapting to changing threat landscapes [11]. These hybrid methods significantly improve prioritisation accuracy, yet challenges remain in ensuring explainability and scalability in high-velocity production environments.

2.2 Knowledge Extraction from Alerts

Beyond ranking individual alerts, extracting actionable knowledge—such as attack strategies, causal relationships, and behavioral patterns—is crucial for understanding the broader threat context. This knowledge is often derived through correlation, graph construction, and pattern mining.

Attribute-similarity methods represent an early and efficient approach to knowledge extraction, correlating alerts based on shared metadata like IP addresses or timestamps to form basic attack narratives [12]. However, these methods are vulnerable to obfuscation techniques. Causality-based methods offer a deeper level of insight by linking alerts through prerequisites and consequences, enabling the reconstruction of logical attack paths as demonstrated in foundational work [13] and extended in stream-mining frameworks like RTECA [14] and deep causal decomposition for industrial systems [15].

Attack graphs provide a structured representation of potential multi-stage attacker behavior. By mapping observed alerts onto these graphs, researchers can model complex campaigns and improve interpretability for analysts [16]. Probabilistic extensions, such as Bayesian attack graphs and trees, incorporate uncertainty and vulnerability severity to enrich the extracted knowledge [17]. More recent work leverages graph-based correlation frameworks [18] and automaton-based learning like S-PDFA [19] to derive compact representations of attack patterns.

Hybrid approaches integrate these techniques with modern AI. The use of large language models (LLMs) for alert management represents a frontier in extracting semantic meaning and providing natural language summaries of attack campaigns [20]. Despite these advances, achieving a balance between the depth of extracted knowledge, computational efficiency for real-time use, and generalization to novel attack sequences remains an open research problem. The interplay between effective prioritisation and deep knowledge extraction is essential for building next-generation SOC systems.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

3 Methodology

In this section, we propose the TTP-Guided Event Reconstruction Framework, which consists of three layers: Semantic Knowledge, Attack Modeling, and Event Reconstruction, as illustrated in Fig. 1.

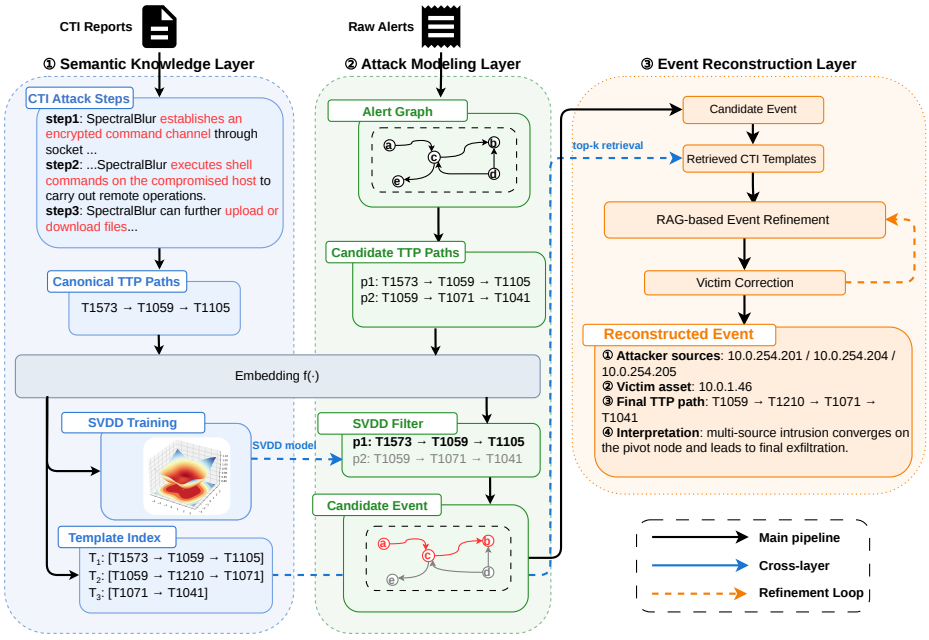

Fig. 1. Overall Architecture of the Proposed TTP-Guided Event Reconstruction Framework

The Semantic Knowledge Layer builds a structured attack knowledge space. It transforms threat intelligence into template-based TTP paths, embeds them into vector representations, and trains a Deep SVDD model to learn the distribution of valid attack patterns. The Attack Modeling Layer constructs candidate attack structures from alerts. It organizes alerts into an attack graph, extracts TTP paths, embeds them into vectors, and applies the trained SVDD model to filter structurally abnormal paths. The remaining paths are aggregated into candidate events. The Event Reconstruction Layer refines these candidate events through retrieval-augmented reasoning. Relevant templates are retrieved from the knowledge space based on TTP path similarity, and an LLM generates structured event diagnostics. Victim correction and convergence checks are applied to ensure consistency.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

3.1 Semantic Knowledge Layer

The purpose of the Semantic Knowledge Layer is to construct a comparable semantic representation of attack behaviors. This representation serves two functions: (1) building a vector retrieval index for template alignment, and (2) learning the structural distribution of valid attack patterns for anomaly filtering.

Given a threat intelligence report, we employ an LLM to reconstruct explicit attack processes under a predefined schema constraint. The model extracts one or multiple attack chains, each represented as an ordered sequence of structured steps:

$$\mathcal{C} = (s_1, s_2, \ldots, s_k) \tag{1}$$

where each step s_i contains attack action, involved entities, temporal information, mapped MITRE ATT&CK tactic/technique identifiers, and supporting evidence.

From each extracted chain $\mathcal{C}$, we construct an attack query graph:

$$G = (V, E) \tag{2}$$

where each node in V corresponds to a step s_i, and each directed edge in E encodes temporal progression according to the original step index. No graph traversal reordering is performed; the structure preserves the chronological sequence extracted from the report.

For structural comparison and retrieval, the query graph G is serialized into a canonical TTP path:

$$P = (t_1, t_2, \ldots, t_n) \tag{3}$$

where t_i denotes the technique identifier of step s_i sorted by step index. Before indexing, the sequence undergoes normalization: unknown or missing techniques are removed, technique identifiers are standardized to the form $T1046$ or $T1059.003$, and sub-techniques not present in the known set are optionally reduced to their base techniques. The resulting normalized sequence is treated as the canonical structural signature of the template.

Templates sharing the same normalized TTP path P are merged into a single template identifier. Each template is therefore defined as:

$$\mathcal{T} = (G, P, M) \tag{4}$$

where M stores aggregated metadata, including supporting instances and textual evidence from multiple reports.

To obtain a comparable semantic representation, each canonical TTP path P is embedded into a fixed-dimensional vector:

$$\mathbf{v} = f(P) \tag{5}$$

where $f(\cdot)$ denotes the embedding function and $\mathbf{v} \in \mathbb{R}^d$. For identical normalized TTP paths, the resulting embedding $\mathbf{v}$ is deterministic and reused.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

The embedding vectors $\mathbf{v}$ serve two essential roles. First, they are used to construct a vector retrieval index:

$$\mathcal{I} = \{(\mathbf{v}_i, M_i)\}_{i=1}^{N} \tag{6}$$

implemented using FAISS with inner-product search over L2-normalized vectors, which is equivalent to cosine similarity. The index supports top-k nearest-neighbor retrieval (default $k = 5$) during downstream reasoning.

Second, the embeddings $\{\mathbf{v}_i\}$ are used to train a Deep SVDD model to characterize the structural distribution of valid attack templates. Since the template library is derived from curated threat intelligence reports, all templates are treated as structurally valid samples. Let $\phi(\mathbf{v})$ denote a neural transformation and $\mathbf{c}$ denote the hypersphere center. The training objective is:

$$\min_{W} \frac{1}{N} \sum_{i=1}^{N} \|\phi(\mathbf{v}_i) - \mathbf{c}\|^2 \tag{7}$$

where W represents network parameters. The learned hypersphere defines a compact region corresponding to coherent attack structures in the embedding space. The decision radius is later determined from the empirical distance distribution of training samples.

Through this process, the Semantic Knowledge Layer establishes a unified semantic space that simultaneously supports similarity-based retrieval and structural distribution learning.

3.2 Attack Modeling Layer

The Attack Modeling Layer transforms normalized intrusion alerts into candidate attack events through interaction modeling, constrained path extraction, and structural filtering.

Given a set of normalized alerts, temporally adjacent records are first aggregated into interaction units. We use the following aggregation key:

$$\begin{aligned} k(a_i) = (&\mathit{src\_ip}, \mathit{dst\_ip}, \mathit{dst\_port}, \\ &\mathit{app\_proto}, \mathit{signature\_id}, \mathit{http\_method}, \mathit{http\_path}). \end{aligned} \tag{8}$$

Alerts with the same key are merged when the time gap between two consecutive records does not exceed δ.

Based on the aggregated alerts, we construct a directed alert graph $G_a = (V_a, E_a)$, where nodes denote host or service entities and edges denote aggregated interactions. Each edge is assigned a MITRE ATT&CK technique identifier by first applying rule-based matching and then using the LLM only when the rule result is unknown. The identified techniques are further mapped into semantic categories, including RECON, CONTROL, C2, and IMPACT, to guide constrained path expansion.

Candidate TTP paths are then extracted from G_a using depth-first search under temporal and semantic constraints. For each path P, we compute its

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

semantic embedding $\mathbf{v} = f(P)$ and evaluate its structural consistency with the trained Deep SVDD model. Specifically, the distance to the learned center is computed as

$$d(\mathbf{v}) = \|\phi(\mathbf{v}) - \mathbf{c}\|, \tag{9}$$

and a path is retained as an inlier if $d(\mathbf{v}) \leq R$, where R is the decision radius estimated from the training distribution.

Finally, the retained inlier paths are aggregated into candidate attack events according to temporal proximity and host-level connectivity.

3.3 Event Reconstruction Layer

While the Attack Modeling Layer enforces structural plausibility through constrained path extraction and filtering, structural consistency alone does not guarantee semantic correctness. The Event Reconstruction Layer therefore introduces retrieval-augmented reasoning to improve attacker–victim coherence beyond structural validation.

Each candidate event consists of multiple TTP paths grouped by structural skeleton and temporal proximity. All paths within the group are merged into a weighted directed event graph $G_e = (V_e, E_e)$, where edges sharing identical (u, v) are aggregated and weighted by their structural support.

To obtain a coherent retrieval query, a representative attacker-to-victim path is synthesized directly from the merged event graph. Given attacker A and victim V, we search for a simple path π from A to V under a maximum depth constraint and score each candidate by cumulative support:

$$\text{Score}(\pi) = \sum_{e \in \pi} w(e) \tag{10}$$

The highest-scoring path is selected, and the dominant TTP on each edge is used to form the representative TTP path $P_{\mathcal{E}}$.

The resulting path is embedded into the shared semantic space as $\mathbf{v}_{\mathcal{E}} = f(P_{\mathcal{E}})$, and similarity retrieval is performed over the template index:

$$\mathcal{R}_k = \text{TopK}_k(\mathbf{v}_{\mathcal{E}}, \mathcal{I}) \tag{11}$$

The LLM receives the structured event evidence together with the retrieved template summaries to refine attacker–victim assignments. If a corrected victim is supported by the event evidence, the event representation is updated as

$$\mathcal{E}' = (A, V', \{P'\}, \Delta t', \mathcal{L}') \tag{12}$$

and reconstruction is repeated until convergence or a predefined iteration bound is reached.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

4 Experiments

In this section, we evaluate the proposed framework from three aspects: overall effectiveness, module contribution, and qualitative correctness. Specifically, we examine whether the framework can reconstruct coherent attack events from noisy intrusion alerts, whether its two key modules, SVDD-based structural filtering and RAG-based semantic refinement, provide measurable improvements, and whether the recovered events are consistent with known attack structure, role assignments, and stage progression in representative scenarios.

Since the datasets used in this study do not provide ground-truth attack chains, the evaluation focuses on structural plausibility, semantic consistency, and end-to-end reconstruction quality rather than supervised accuracy alone. Accordingly, we formulate the following research questions (RQs):

- **RQ1:** How effective is the proposed framework in reducing alert redundancy and reconstructing coherent attack events from noisy intrusion alerts? (Section 4.2)
- **RQ2:** What is the contribution of the two key modules in the framework, namely SVDD-based structural filtering and RAG-based semantic refinement? (Section 4.3)
- **RQ3:** Can the reconstructed attack event correctly reflect the known attack structure, role assignments, and stage progression in a representative real-world attack scenario? (Section 4.4)

4.1 Experimental Setup

Datasets We evaluate the proposed framework on three publicly available intrusion alert datasets, namely CPTC-2017, CPTC-2018, and CCDC-2018. These datasets are collected from realistic cyber range exercises and contain multi-stage intrusion behaviors. All samples are provided in the form of Suricata alerts, including timestamps, source and destination IP addresses and ports, and alert signature identifiers. Table 1 summarizes the basic statistics of the adopted datasets.

Table 1. Experimental Dataset Summary

Property	CPTC-2017	CPTC-2018	CCDC-2018
# **alerts**	43,611	330,270	1,052,281
# **teams**	9	6	–
# **IPs**	494	42	2,138
Duration (hours)	11	9	25
Dataset type	Penetration Testing	Penetration Testing	Blue-Team Testing

Since these datasets do not provide ground-truth attack chains, our evaluation does not rely on direct chain-matching accuracy. Instead, we focus on structural plausibility, noise reduction capability, and semantic consistency.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

All experiments follow the three-layer framework described in Section 3. The Semantic Knowledge Layer is trained using attack templates extracted from publicly available cyber threat intelligence reports describing real-world attack campaigns. The CTI corpus is independent of the evaluated alert datasets, ensuring that no evaluation alerts are involved in template construction or SVDD training.

Canonical TTP paths derived from the CTI templates are embedded using the `text-embedding-3-small` model and used to train a Deep SVDD model. The decision radius R is determined during training as the $(1-\nu)$-quantile of the training distance distribution, where $\nu = 0.2$ in all experiments.

For alert-side processing, raw alerts are first aggregated using a time-window parameter of $\delta = 600$ seconds. Aggregated interactions are then mapped to MITRE ATT&CK techniques through rule-based inference, with LLM assistance applied only when rule matching fails. Candidate TTP paths are extracted using constrained depth-first search with a maximum hop limit of 5.

All candidate paths are embedded using the same embedding function $f(\cdot)$ and evaluated by the trained SVDD model. Structurally consistent paths are further aggregated into candidate events. Similarity retrieval is implemented with FAISS using inner-product search over L2-normalized vectors, which is equivalent to cosine similarity. For each event, the top-k nearest templates are retrieved with $k = 5$. Unless otherwise stated, all parameters are fixed across datasets.

Metrics Because no ground-truth attack chains are available, we evaluate the proposed framework using indirect metrics that reflect structural filtering quality, event-level coherence, and semantic refinement effectiveness.

To measure the effect of structural filtering, we report the structural reduction rate (SR):

$$SR = 1 - \frac{|P^*|}{|P|}, \tag{13}$$

where $|P|$ denotes the number of candidate paths before SVDD filtering and $|P^*|$ denotes the number of retained paths after filtering.

To characterize the continuity of reconstructed paths, we report the average path length (APL):

$$APL = \frac{1}{|P^*|} \sum_{i=1}^{|P^*|} |C_i|, \tag{14}$$

where $|C_i|$ denotes the number of stages in the i-th retained path. A larger APL indicates that longer multi-stage attack progressions are preserved after filtering.

We also report the TTP retention ratio (TR):

$$TR = \frac{|T_{final}|}{|T_{candidate}|}, \tag{15}$$

where $|T_{candidate}|$ and $|T_{final}|$ denote the numbers of TTP instances before and after structural filtering, respectively.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

At the event level, we evaluate attacker–victim coherence using role consistency and stage integration using stage coverage. Role consistency measures the proportion of reconstructed events whose attacker–victim assignments remain semantically coherent after event-level fusion. Stage coverage measures the proportion of reconstructed events that preserve the expected multi-stage attack progression, especially the joint presence of key stages such as CONTROL, C2, and IMPACT.

For semantic refinement, we report the victim correction rate (VCR), which measures the proportion of reconstructed events whose victim assignment is corrected after RAG-based refinement.

In the ablation study, we additionally report the number of retained paths $|P^*|$ and the number of reconstructed events $|E|$ to show how each module affects the scale of the remaining candidate space and the granularity of the final event reconstruction.

4.2 Quantitative Evaluation (RQ1)

To answer **RQ1**, we present quantitative results on the effectiveness of the proposed framework in suppressing noisy candidate paths and reconstructing coherent attack events.

Table 2 shows that structural filtering substantially reduces candidate TTP paths in CPTC-2018 and CCDC-2018, with structural reduction rates of 79.85% and 55.48% respectively. This indicates that a large portion of the extracted paths do not conform to structurally plausible attack patterns and can be effectively removed before event reconstruction. Although stage-wise reductions are observed after filtering, CONTROL, C2, and IMPACT remain detectable across all datasets, suggesting that the proposed constraints mainly suppress weakly supported paths rather than systematically removing critical attack evidence.

Table 2. Candidate path reduction after structural filtering.

Dataset	Before	After	Reduction (%)
CPTC-2017	484	264	54.54
CPTC-2018	268	54	79.85
CCDC-2018	2790	1242	55.48

Table 3 further shows that event-level fusion consistently improves attacker-victim role consistency across all datasets, with the most substantial gain observed on CCDC-2018. Its effect on stage coverage is more dataset-dependent: fusion improves stage integration in CPTC-2018, whereas the stage coverage of CPTC-2017 decreases slightly after fusion due to the removal of weak or isolated evidence. In CCDC-2018, complete CONTROL-C2-IMPACT coverage remains absent both before and after fusion, indicating that attack evidence is more strongly fragmented across paths in this dataset. Overall, these results show

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

that the proposed framework effectively reduces alert-induced path noise while improving the coherence of reconstructed attack events.

Table 3. Role consistency and stage coverage before and after event-level fusion (%).

Dataset	Role Consistency		Stage Coverage	
	Before	After	Before	After
CPTC-2017	18.6	**46.3**	12.5	6.4
CPTC-2018	57.4	**83.3**	18.5	**25.0**
CCDC-2018	25.1	**97.0**	0.0	0.0

4.3 Ablation Study (RQ2)

To answer **RQ2**, we ablate the two major components of the proposed framework, namely the SVDD-based structural filtering module and the RAG-based semantic refinement module. Since the two components operate at different stages of the pipeline, we analyze their contributions from both the structural and event-level perspectives. Table 4 reports the ablation results of different variants across the three datasets. The symbol "–" denotes either an unchanged value with respect to the Full model or a metric that is not applicable in that setting.

Table 4. Ablation results of the proposed framework.

Dataset	Variant	$\lvert P^* \rvert$	$\lvert E \rvert$	APL	TR	VCR
CPTC-2017	Full	264	188	3.79	0.63	0.60
	w/o SVDD	484	339	3.27	1.00	–
	w/o RAG	–	188	–	–	0.00
CPTC-2018	Full	54	12	3.41	0.24	0.58
	w/o SVDD	268	116	2.87	1.00	–
	w/o RAG	–	26	–	–	0.00
CCDC-2018	Full	1242	165	3.55	0.54	0.38
	w/o SVDD	2790	1033	2.94	1.00	–
	w/o RAG	–	606	–	–	0.00

Table 4 shows that removing SVDD substantially enlarges the retained path set and increases the number of reconstructed events across all datasets. It also reduces APL, indicating that the SVDD-based filtering module suppresses structurally weak path combinations while preserving more coherent multi-stage attack progressions. This effect is especially pronounced on CCDC-2018, where the candidate space is much larger and structural control becomes more important.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

In contrast, removing RAG leaves the structural metrics unchanged but increases the number of fragmented events and reduces VCR to zero. This indicates that the RAG-based module does not alter the underlying path topology, but instead improves event consolidation and attacker–victim assignment at the semantic level.

Overall, the ablation results show that the two modules contribute in complementary ways. SVDD improves structural quality by constraining the candidate search space, whereas RAG improves semantic quality by consolidating event representations and refining attacker–victim assignments.

4.4 Case Study (RQ3)

To answer **RQ3**, we present a representative case from the CPTC-2018 dataset to examine whether the reconstructed event is consistent with the known attack background and the expected attacker-victim structure.

According to the public competition description and the S-PDFA report, hosts 10.0.254.201, 10.0.254.204, and 10.0.254.205 correspond to different red-team attack sources, while 10.0.1.46 is the victim asset. As shown in Fig. 2, the reconstructed result exhibits a clear multi-source converging structure: different attack sources establish interactions with node 10.0.0.24 through execution and exploitation behaviors (T1059 and T1210) as well as command-and-control communication (T1071), and the intermediate node then transfers data to the victim asset 10.0.1.46 via T1041. This recovered structure is consistent with the expected progression from parallel intrusion attempts to core-node convergence and final exfiltration.

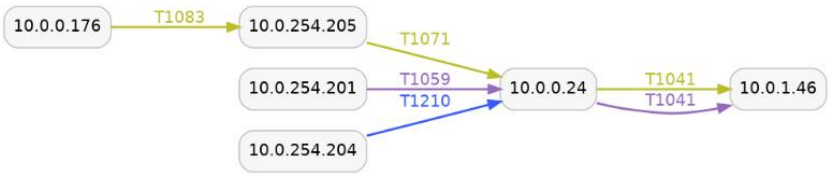

Fig. 2. A representative multi-source converging attack event reconstructed from CPTC-2018. The recovered structure is consistent with the known attacker-victim roles and stage progression in the exercise scenario.

At the intermediate stages of reconstruction, the corresponding alert region contains multiple heterogeneous alerts generated by different source hosts. After alert aggregation, these alerts are condensed into interaction-level edges with mapped TTP evidence. Candidate path generation further enumerates several topologically reachable path variants, including short partial paths and structurally incomplete combinations. After SVDD-based filtering, only those paths that preserve plausible stage continuity are retained, which suppresses fragmented variants and preserves the main four-step attack progression.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

Finally, event-level reconstruction consolidates the retained paths into a unified attack event centered on node 10.0.0.24. In the final result, the attacker-side sources, the intermediate compromised node, and the downstream victim are organized into a coherent attacker-victim structure that matches the known attack background. Therefore, this case suggests that the proposed framework does not merely produce a plausible event representation, but can recover an attack structure that is consistent with the actual exercise scenario.

Overall, this case study provides qualitative evidence that the proposed framework can recover attack-event structure, role assignments, and stage progression that are consistent with a representative multi-source attack scenario.

5 Conclusion

In this paper, we proposed a TTP-Guided Event Reconstruction Framework for reconstructing coherent attack events from large volumes of intrusion alerts. By introducing CTI-derived canonical TTP paths into the reconstruction process, the framework connects threat knowledge with alert-side interaction evidence and supports structure-aware attack-event recovery.

Experiments on three public cyber range datasets show that the proposed framework can effectively suppress noisy candidate paths and improve attacker–victim coherence in reconstructed events. The ablation study further demonstrates that the two key modules play complementary roles: SVDD-based filtering improves structural quality, while RAG-based refinement enhances event consolidation and role assignment. In addition, the case study shows that the reconstructed event is consistent with the known attack background and stage progression in a representative attack scenario.

Although the current evaluation is limited by the absence of ground-truth attack chains in public datasets, the results suggest that TTP sequences provide an effective semantic bridge between CTI knowledge and alert observations. In future work, we will further evaluate the framework on datasets with more explicit attack-chain annotations.

Acknowledgments. This work is supported by the National Key R&D Program of China (No. 2024YFB3109004), Natural Science Foundation of China under Grant No. 62202466 and Youth Innovation Promotion Association CAS under Grant No. 2022159. This research was also supported by Key Laboratory of Network Assessment Technology, Chinese Academy of Sciences, and Beijing Key Laboratory of Network Security and Protection Technology.

References

1. Faris Bugra Kokulu, Ananta Soneji, Tiffany Bao, Yan Shoshitaishvili, Ziming Zhao, Adam Doupé, and Gail-Joon Ahn. Matched and mismatched socs: A qualitative study on security operations center issues. In *Proceedings of the 2019 ACM SIGSAC conference on computer and communications security*, pages 1955–1970, 2019.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

2. Jun Zengy, Xiang Wang, Jiahao Liu, Yinfang Chen, Zhenkai Liang, Tat-Seng Chua, and Zheng Leong Chua. Shadewatcher: Recommendation-guided cyber threat analysis using system audit records. *2022 IEEE Symposium on Security and Privacy (SP)*, pages 489–506, 2022.
3. Xander Bouwman, Harm Griffioen, Jelle Egbers, Christian Doerr, Bram Klievink, and Michel Van Eeten. A different cup of {TI}? the added value of commercial threat intelligence. In *29th USENIX security symposium (USENIX security 20)*, pages 433–450, 2020.
4. Yuxuan Jiang, Chaoyun Zhang, Shilin He, Zhihao Yang, Ming-Jie Ma, Si Qin, Yu Kang, Yingnong Dang, S. Rajmohan, Qingwei Lin, and Dongmei Zhang. Xpert: Empowering incident management with query recommendations via large language models. *2024 IEEE/ACM 46th International Conference on Software Engineering (ICSE)*, pages 1121–1133, 2023.
5. Thijs van Ede, Hojjat Aghakhani, Noah Spahn, Riccardo Bortolameotti, Marco Cova, Andrea Continella, Maarten van Steen, Andreas Peter, Christopher Kruegel, and Giovanni Vigna. DEEPCASE: semi-supervised contextual analysis of security events. In *43rd IEEE Symposium on Security and Privacy, SP 2022*, pages 522–539. IEEE, 2022.
6. Chuanpu Fu, Qi Li, Ke Xu, and Jianping Wu. Point cloud analysis for ml-based malicious traffic detection: Reducing majorities of false positive alarms. In *Proceedings of the 2023 ACM SIGSAC Conference on Computer and Communications Security*, pages 1005–1019, 2023.
7. Max Landauer, Florian Skopik, Markus Wurzenberger, and Andreas Rauber. Dealing with security alert flooding: Using machine learning for domain-independent alert aggregation. *ACM Trans. Priv. Secur.*, 25(3):18:1–18:36, 2022.
8. Fatemeh Jalalvand, Mohan Baruwal Chhetri, Surya Nepal, and Cecile Paris. Alert prioritisation in security operations centres: A systematic survey on criteria and methods. *ACM Computing Surveys*, 57(2):1–36, 2024.
9. Jianguo Jiang, Qiwen Wang, Zhixin Shi, Bin Lv, Wei Fan, and Xiao Peng. The parameter optimization based on lvpso algorithm for detecting multi-step attacks. In *Proceedings of the 16th ACM International Conference on Computing Frontiers*, pages 24–31, 2019.
10. Peng Zhou, Gongyan Zhou, Dakui Wu, and Minrui Fei. Detecting multi-stage attacks using sequence-to-sequence model. *Computers & Security*, 105:102203, 2021.
11. Xiaoyu Wang, Xiaobo Yang, Xueping Liang, Xiu Zhang, Wei Zhang, and Xiaorui Gong. Combating alert fatigue with alertpro: Context-aware alert prioritization using reinforcement learning for multi-step attack detection. *Computers & Security*, 137:103583, 2024.
12. Xiang Cheng, Jiale Zhang, and Bing Chen. Cyber situation comprehension for iot systems based on apt alerts and logs correlation. *Sensors*, 19(18):4045, 2019.
13. Peng Ning, Yun Cui, and Douglas S Reeves. Constructing attack scenarios through correlation of intrusion alerts. In *Proceedings of the 9th ACM Conference on Computer and Communications Security*, pages 245–254, 2002.
14. Ali Ahmadian Ramaki, Morteza Amini, and Reza Ebrahimi Atani. Rteca: Real time episode correlation algorithm for multi-step attack scenarios detection. *computers & security*, 49:206–219, 2015.
15. Zahra Jadidi, Joshua Hagemann, and Daniel Quevedo. Multi-step attack detection in industrial control systems using causal analysis. *Computers in Industry*, 142:103741, 2022.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18

16. Hao Hu, Jing Liu, Yuchen Zhang, Yuling Liu, Xiaoyu Xu, and Jinglei Tan. Attack scenario reconstruction approach using attack graph and alert data mining. *Journal of Information Security and Applications*, 54:102522, 2020.
17. Xinzhou Qin and Wenke Lee. Attack plan recognition and prediction using causal networks. In *20th Annual Computer Security Applications Conference*, pages 370–379. IEEE, 2004.
18. Yanbang Wang, Karl Hallgren, and Jonathan Larson. A graph-based framework for reducing false positives in authentication alerts in security systems. In *Companion Proceedings of the ACM Web Conference 2024*, pages 274–283, 2024.
19. Azqa Nadeem, Sicco Verwer, Stephen Moskal, and Shanchieh Jay Yang. Alert-driven attack graph generation using s-pdfa. *IEEE transactions on dependable and secure computing*, 19(2):731–746, 2021.
20. Daniel Adanza, LLuis Gifre, Pol Alemany, Carlos Natalino, Paolo Monti, Raul Muñoz, and Ricard Vilalta. Leveraging generative ai for intent-based networking operations in network slices. *Computer Networks*, page 111647, 2025.
21. Min Du, Feifei Li, Guineng Zheng, and Vivek Srikumar. Deeplog: Anomaly detection and diagnosis from system logs through deep learning. In *Proceedings of the 2017 ACM SIGSAC Conference on Computer and Communications Security, CCS 2017*, pages 1285–1298. ACM, 2017.
22. Dongqi Han, Zhiliang Wang, Wenqi Chen, Ying Zhong, Su Wang, Han Zhang, Jiahai Yang, Xingang Shi, and Xia Yin. Deepaid: Interpreting and improving deep learning-based anomaly detection in security applications. In *CCS '21: 2021 ACM SIGSAC Conference on Computer and Communications Security*, pages 3197–3217. ACM, 2021.
23. Limin Yang, Wenbo Guo, Qingying Hao, Arridhana Ciptadi, Ali Ahmadzadeh, Xinyu Xing, and Gang Wang. CADE: detecting and explaining concept drift samples for security applications. In *30th USENIX Security Symposium, USENIX Security 2021*, pages 2327–2344. USENIX Association, 2021.
24. Marie-Andrée Gardella, Pablo Musé, Jean-Michel Morel, and Miguel Colom. Noisesniffer: a fully automatic image forgery detector based on noise analysis. *2021 IEEE International Workshop on Biometrics and Forensics (IWBF)*, pages 1–6, 2021.
25. Sanjana Ingale, Milind Paraye, and Dayanand Ambawade. A survey on methodologies for multi-step attack prediction. In *2020 Fourth International Conference on Inventive Systems and Control (ICISC)*, pages 37–45, 2020.
26. CyberMonitor. Apt cybercriminal campagin collections, 2024.
27. Pilar Holgado, Víctor A Villagrá, and Luis Vazquez. Real-time multistep attack prediction based on hidden markov models. *IEEE Transactions on Dependable and Secure Computing*, 17(1):134–147, 2017.
28. Timothy Chadza, Konstantinos G Kyriakopoulos, and Sangarapillai Lambotharan. Analysis of hidden markov model learning algorithms for the detection and prediction of multi-stage network attacks. *Future generation computer systems*, 108:636–649, 2020.
29. Zhiyong Luo, Xu Yang, Jiahui Liu, and Rui Xu. Network intrusion intention analysis model based on bayesian attack graph. *Journal on Communications*, 41(9):160–169, 2020.
30. Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. Deep one-class classification. In *International conference on machine learning*, pages 4393–4402. PMLR, 2018.

ICCS Camera Ready Version 2026
To cite this paper please use the final published version:
DOI: 10.1007/978-3-032-29921-5_18