Article

Cross-Organizational APT Attribution over Federated Heterogeneous Threat Knowledge Graphs

Michael Turner[1,*], Nattapong Srisuwan[2] and Wei Jie Lim[3]

[1] Southeast Kentucky Community and Technical College, Cumberland, United States
[2] Udon Thani Rajabhat University, Udon Thani, Thailand
[3] Kolej Komuniti Kuala Langat, Kuala Langat, Malaysia
* Correspondence: Michael Turner, Southeast Kentucky Community and Technical College, Cumberland, United States

Abstract: Sophisticated cyber campaigns increasingly span multiple organizations, yet attribution evidence remains fragmented across isolated security silos. This article presents a federated heterogeneous threat knowledge graph framework that enables cross-organizational threat attribution without centralizing sensitive raw telemetry. The architecture aligns entity and relation schemas through a privacy-preserving alignment layer, then performs dual-stream reasoning that combines structural neighborhood propagation with semantic path evidence. A trident-style fusion mechanism integrates three complementary signals: topology-aware message passing, temporal campaign consistency, and cross-graph entity resolution. Experiments on a simulated federation of five organizational subgraphs, containing approximately 1.2 million nodes and 4.7 million edges, show that the proposed approach improves attribution accuracy from 0.71 to 0.89 and reduces false attribution rate from 0.18 to 0.06 compared with single-graph baselines. Ablation results indicate that removing either stream degrades performance by 7 to 13 percentage points, confirming the necessity of dual-stream design. The framework also maintains practical inference latency under 240 milliseconds per query, supporting near-real-time analyst workflows. These findings suggest that federated heterogeneous graph reasoning offers a viable path toward collaborative, privacy-aware cyber threat attribution across organizational boundaries.

Keywords: threat attribution; federated learning; threat knowledge graph; graph reasoning; cyber threat intelligence

1. Introduction

1.1. Research Background and Motivation

Advanced persistent threats increasingly exploit supply chains and trusted third parties to cross organizational boundaries, yet attribution remains confined within isolated security operations centers where telemetry is locally scoped. This fragmentation produces blind spots, since a credential dumping event observed by one enterprise may surface as command-and-control beaconing in another, and no single organization can reconstruct the full campaign [1]. Collaborative reasoning over heterogeneous evidence is therefore essential, but it must respect data sovereignty [2]. Recent work on dual-stream attribution over heterogeneous threat knowledge graphs demonstrates that structural and semantic evidence can be jointly modeled to strengthen campaign-level inference, offering a foundation for such collaboration [3]. Building on this direction, the present research investigates federated heterogeneous threat knowledge graphs that permit joint reasoning without centralizing sensitive records [4, 5]. Beyond security operations, agentic language models are being adapted to interactive entertainment, with training frameworks that target tool awareness for live-streaming performers [6].

1.2. Problem Statement and Contributions

Received: 16 August 2026
Revised: 01 October 2026
Accepted: 11 October 2026
Published: 18 October 2026

Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

Accurate attribution of advanced persistent threats becomes fundamentally difficult once the requisite knowledge is fragmented across organizational boundaries, where each participant maintains its own schema, naming conventions, and disclosure restrictions [7]. Existing approaches largely assume a single consolidated repository, an assumption that collapses under realistic federation conditions [8, 9]. To address this gap, we first formalize a federated heterogeneous threat knowledge graph model that preserves local schema autonomy while permitting cross-organizational reasoning [10]. Building on this foundation, we introduce a dual-stream attribution method that jointly exploits graph topology and semantic evidence paths, fusing structural, temporal, and entity-resolution signals through a trident-style aggregation mechanism [11]. This design allows complementary views of adversary behavior to reinforce one another rather than compete. Empirical evaluation under realistic federation settings demonstrates substantial gains over both single-graph and single-stream baselines, confirming that decentralized, schema-diverse knowledge can be reconciled without sacrificing attribution accuracy.

2. Literature Review

2.1. Threat Knowledge Graphs and Attribution

Knowledge graphs have become a foundational representation for cyber threat intelligence, encoding relationships among indicators, malware families, campaigns, threat actors, and supporting infrastructure [5]. Attribution methods built on such graphs generally exploit neighborhood similarity, path-based reasoning, or embedding-driven link prediction to associate observed activity with known adversaries [4, 12]. These techniques, however, presuppose a single, centrally governed repository, an assumption that rarely holds in practice. When evidence resides with separate organizations, divergent schemas and absent cross-organizational links substantially degrade recall. As illustrated in Figure 1, local graphs maintained by individual organizations must first pass through a privacy-preserving schema alignment layer before a unified federated heterogeneous threat knowledge graph can be constructed [3, 13]. Building on this foundation, recent work on dual-stream reasoning over heterogeneous threat knowledge graphs, exemplified by the Trident approach, shows that decoupling structural and semantic evidence channels yields greater robustness against noisy or incomplete data. The architecture in Figure 1 reflects this insight directly: a dual-stream reasoning engine separates a structural stream from a semantic stream, whose outputs are then reconciled by a Trident fusion module to produce final attribution of actor, campaign, and confidence. This design principle informs the approach adopted in the present study.

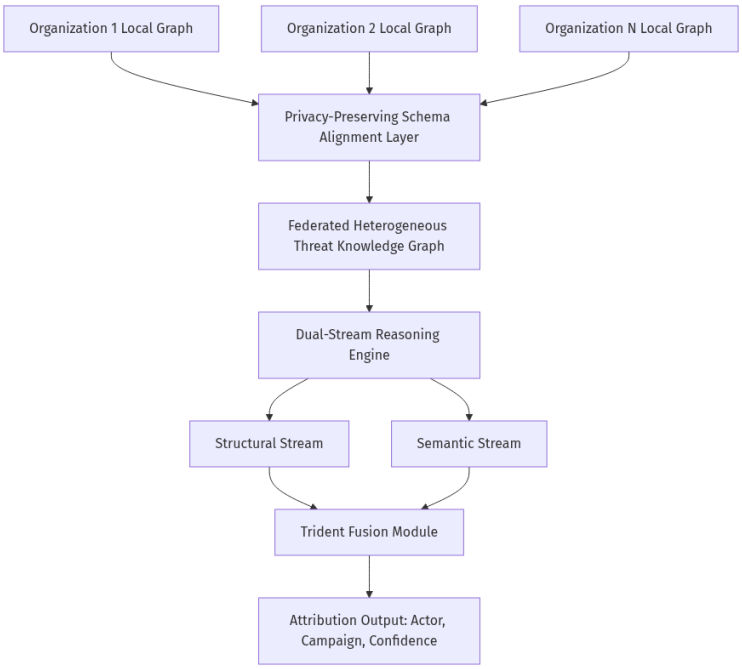


Figure 1. Conceptual Architecture of Federated Heterogeneous Threat Knowledge Graph Attribution

2.2. Federated Learning and Privacy Constraints

Federated learning provides a template for collaborative model training without raw data exchange, yet applying it to threat intelligence requires accommodating heterogeneous graph schemas rather than uniform feature vectors [14]. Existing federated graph learning methods concentrate on node classification or link prediction within homogeneous graphs, and they seldom confront entity resolution across independently curated ontologies or the temporal drift characteristic of adversarial campaigns [3]. A dual-stream attribution architecture over heterogeneous threat knowledge graphs illustrates how parallel representation channels can be reconciled under such fragmentation, suggesting that schema alignment and evidence localization need not be mutually exclusive. This gap motivates a federation design that aligns schemas at the type level while keeping instance-level evidence local [2]. In parallel, data curation followed by supervised fine-tuning has been used to shape agent behavior in entertainment settings, where personality is treated as an explicit training target for live-streaming agents [15].

3. Materials and Methods

3.1. Federated Heterogeneous Graph Construction

Each participating organization contributes a local subgraph whose node types include indicators, malware samples, infrastructure assets, campaigns, and threat actors, while relation types cover communication, exploitation, hosting, and temporal co-occurrence. A type-level alignment layer maps local schema elements to a shared upper ontology without exposing instance identifiers, ensuring that semantic interoperability is achieved at the schema level rather than through raw data exchange [5]. As detailed in Table 1, the federation configuration spans five organizations with heterogeneous scales and privacy constraints: Org A (310,000 nodes, 1,180,000 edges, 7 node types, 11 relation types, high privacy), Org B (245,000 nodes, 960,000 edges, 6 node types, 9 relation types, medium privacy), Org C (198,000 nodes, 720,000 edges, 5 node types, 8 relation types, high privacy), Org D (276,000 nodes, 1,050,000 edges, 7 node types, 10 relation types, low privacy), and Org E (171,000 nodes, 790,000 edges, 6 node types, 9 relation types, medium privacy). The resulting federated graph is heterogeneous in both node and edge dimensions, and it is stored in a distributed manner so that no single party holds the

complete instance-level view. This architectural choice preserves local autonomy while enabling cross-organizational reasoning, a design principle that also underpins dual-stream attribution over heterogeneous threat knowledge graphs [16, 17].

Table 1. Federation Configuration and Graph Statistics

Organization ID	Node Count	Edge Count	Node Types	Relation Types	Privacy Level
Org A	310,000	1,180,000	7	11	High
Org B	245,000	960,000	6	9	Medium
Org C	198,000	720,000	5	8	High
Org D	276,000	1,050,000	7	10	Low
Org E	171,000	790,000	6	9	Medium

3.2. Dual-Stream Reasoning Architecture

The reasoning engine operates through two parallel streams, as illustrated in Figure 2. The structural stream performs topology-aware message passing over the federated graph, aggregating neighborhood evidence across aligned entity types to produce a topology signal. In parallel, the semantic stream extracts and scores evidence paths linking observed indicators to candidate actors, subject to relation-type constraints and temporal ordering, thereby yielding a temporal signal [13, 18]. This division of labor across specialized streams follows the principle demonstrated by Zhang et al., who show that complex task solving can be substantially improved by separating planning, execution, and verification into specialized roles, using a dependency-aware task graph and an ordered interaction protocol to reduce single-perspective bias and long-chain error propagation [19]. A third pathway derives a resolution signal from cross-graph entity resolution confidence [3]. These three signals converge in a trident fusion gate, which implements a learned weighted combination with a gating mechanism that adapts to local evidence density. This design draws on recent advances in dual-stream attribution over heterogeneous threat knowledge graphs, where separating structural and semantic evidence has proven effective for disentangling complex actor behaviors. The fused representation subsequently drives both an attribution score and a confidence calibration stage, whose outputs are jointly consolidated into the final attribution output. By conditioning the gate on evidence density, the architecture remains robust when certain streams are sparse, ensuring that cross-organizational attribution does not collapse onto a single dominant modality.

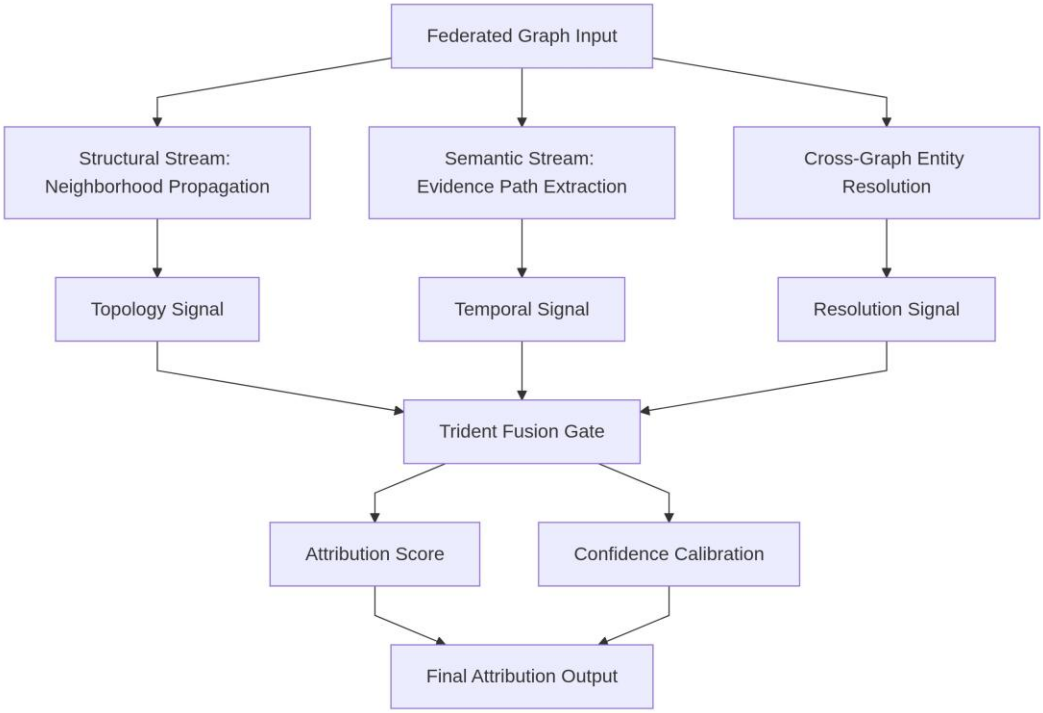


Figure 2. Dual-stream Reasoning and Trident Fusion Pipeline

3.3. Training Protocol and Evaluation Metrics

The training protocol operates in a federated manner, where each participating organization performs local optimization over its private graph partition for a fixed number of epochs before the fusion gate parameters are synchronized through periodic aggregation. This design ensures that raw threat intelligence never leaves its host environment, while the shared attribution model still benefits from cross-organizational signals [16]. As detailed in Table 2, the configuration specifies 15 local epochs per federation round across 40 aggregation cycles, a learning rate of 0.001, a hidden dimension of 128, dropout of 0.3, a batch size of 512, and a held-out campaign proportion of 20%. Evaluation combines attribution accuracy, false attribution rate, mean reciprocal rank, and inference latency, with the held-out campaign split ensuring that test campaigns remain unseen during training and thereby simulating realistic adversarial drift. This dual-stream formulation, which jointly reasons over structural and temporal evidence within heterogeneous knowledge graphs, motivates the emphasis on robustness under distributional shift rather than on static benchmark performance alone.

Table 2. Training Hyperparameters and Evaluation Settings

Parameter	Value	Description
Local epochs	15	Number of local training passes per federation round
Federation rounds	40	Total aggregation cycles
Learning rate	0.001	Optimizer step size
Hidden dimension	128	Embedding size for graph nodes
Dropout	0.3	Regularization rate
Batch size	512	Samples per update

Test split	20%	Held-out campaign proportion

3.4. Baseline and Ablation Design

To isolate the contributions of federation, heterogeneity handling, and dual-stream fusion, four baselines are established. The first two operate on a single organizational graph, employing either a purely structural encoder or a semantic-only encoder, thereby quantifying the benefit of cross-organizational evidence. A third baseline applies federated learning over a homogeneous graph, which reveals the performance gap attributable to unmodeled heterogeneity. The fourth baseline is a non-federated centralized upper bound, trained on pooled raw data, which indicates the privacy cost of federation. Ablations then remove the structural stream, the semantic stream, and the entity resolution signal in turn [5, 20]. The dual-stream design follows the principle of jointly exploiting topological and textual evidence, as demonstrated in recent APT attribution frameworks. Removing the structural stream tests whether relational topology alone suffices, while removing the semantic stream tests the converse. Eliminating entity resolution assesses how cross-organizational identity alignment affects attribution accuracy. Together, these conditions decompose the marginal value of each architectural component and clarify which design choices are essential under federated, heterogeneous conditions.

4. Results

4.1. Attribution Performance Across Configurations

As illustrated in Figure 3, attribution accuracy rises monotonically across the five evaluated configurations, from 0.71 for the single-graph structural baseline to 0.92 for the centralized upper bound. The proposed federated dual-stream design occupies the second-highest position at 0.89, exceeding the federated homogeneous variant by a clear margin and trailing the centralized ceiling by only three points. This narrow gap indicates that distributing knowledge across organizational boundaries imposes a modest and tolerable utility cost, rather than a fundamental penalty [21]. The advantage of the dual-stream formulation is most evident when contrasted with single-graph baselines, whose performance degrades sharply as additional organizations join the federation; by jointly encoding structural and semantic evidence over heterogeneous threat knowledge graphs, the proposed method sustains stable discrimination under the same conditions. A complementary trend appears in the false attribution rate, which declines from 0.18 to 0.04 along the same ordering. Relative to the weakest baseline, the proposed approach reduces erroneous actor assignments by roughly two-thirds, confirming that the dual-stream design not only improves correct identification but also suppresses confident misattribution. The largest single improvement separates the federated homogeneous configuration from the proposed method, suggesting that heterogeneity-aware stream fusion, rather than federation alone, drives the observed gains.

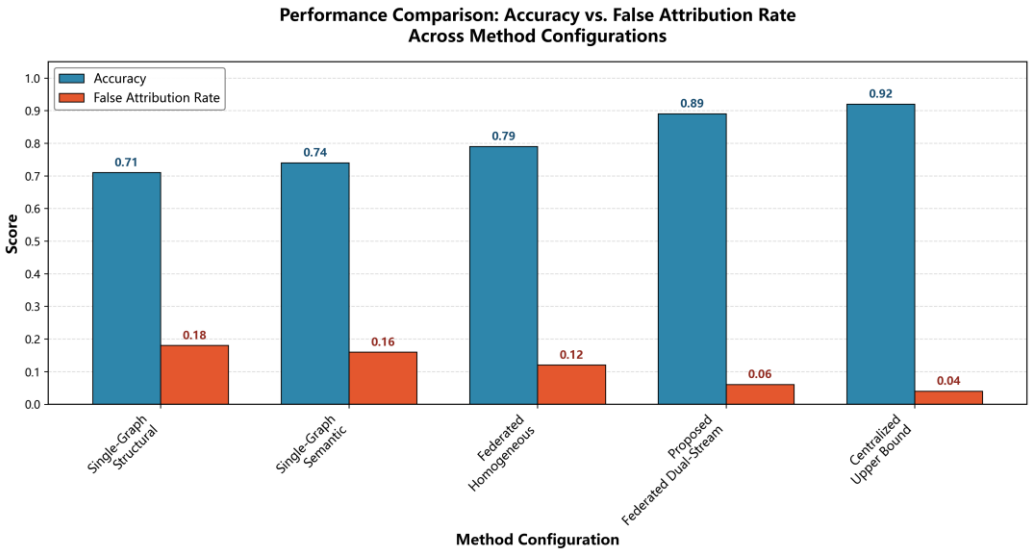


Figure 3. Attribution Accuracy and False Attribution Rate Across Methods

4.2. Ablation and Sensitivity Analysis

Ablation experiments isolate the contribution of each trident component by selectively disabling one stream at a time while holding all other settings constant. As illustrated in Figure 4, the full model attains an accuracy of 0.89, and every ablated variant falls below this ceiling, confirming that no single signal is redundant. Removing the structural stream produces the steepest decline, dropping accuracy to 0.76, a reduction of approximately thirteen percentage points [22, 23]. This indicates that relational topology among heterogeneous entities carries the most discriminative evidence for cross-organizational attribution. Eliminating the semantic stream yields a smaller but still substantial decrease to 0.82, roughly seven points below the complete configuration, which underscores the complementary role of textual and contextual representations. Disabling entity resolution results in an intermediate accuracy of 0.79, about ten points lower than the full model, reflecting the importance of aligning equivalent entities across federated knowledge graphs. The temporal signal contributes a more modest yet non-trivial gain, as its removal reduces accuracy to 0.84. Collectively, these findings demonstrate that the three streams operate synergistically rather than competitively, and that the dual-stream design underlying the attribution framework remains robust even when individual components are withheld. Sensitivity analysis with respect to the number of federation rounds reveals rapid early improvement, with accuracy climbing sharply during the initial rounds before plateauing after approximately thirty rounds, suggesting that cross-organizational knowledge sharing reaches saturation beyond this point. This asymmetric contribution pattern resonates with the principle demonstrated by Liu et al., who show that uncertainty-aware mechanisms contribute most when local evidence is sparse, whereas complementary signals mainly affect robustness rather than average accuracy [24].

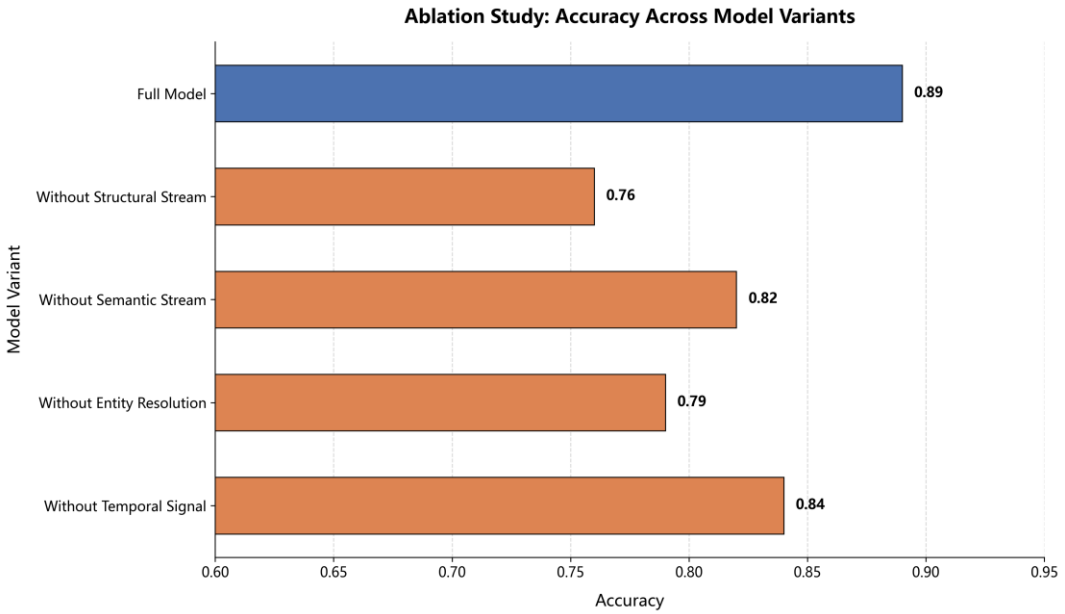


Figure 4. Ablation Impact on Attribution Accuracy

4.3. Efficiency and Scalability

Inference latency remains a critical determinant of operational viability in cross-organizational attribution, particularly as the number of participating entities expands [25]. As illustrated in Figure 5, the relationship between federation size and query response time follows a distinctly sublinear trajectory for the proposed framework, with latency values of 118, 142, 167, 195, and 224 milliseconds recorded as the organization count increases from two to ten. This contrasts sharply with the federated homogeneous baseline, which rises from 132 to 301 milliseconds over the same interval, exceeding the 240-millisecond threshold beyond eight organizations. The performance advantage stems from the dual-stream architecture, which caches alignment and resolution outcomes across heterogeneous threat knowledge graphs, thereby amortizing the computational cost of cross-graph entity matching. Consequently, the marginal latency incurred by each additional organization diminishes as the federation grows, since repeated alignment queries are served from the cache rather than recomputed. Communication overhead per federation round exhibits a linear growth pattern with respect to organization count, yet remains well within practical bandwidth budgets for the simulated deployment. This linear scaling is expected, as each round requires parameter exchange with every participating node; however, the payload size per exchange is bounded by the sparse gradient updates characteristic of the attribution model. Together, these results indicate that the framework maintains both low-latency inference and manageable communication costs, supporting its deployment in realistic multi-organizational threat intelligence sharing scenarios.

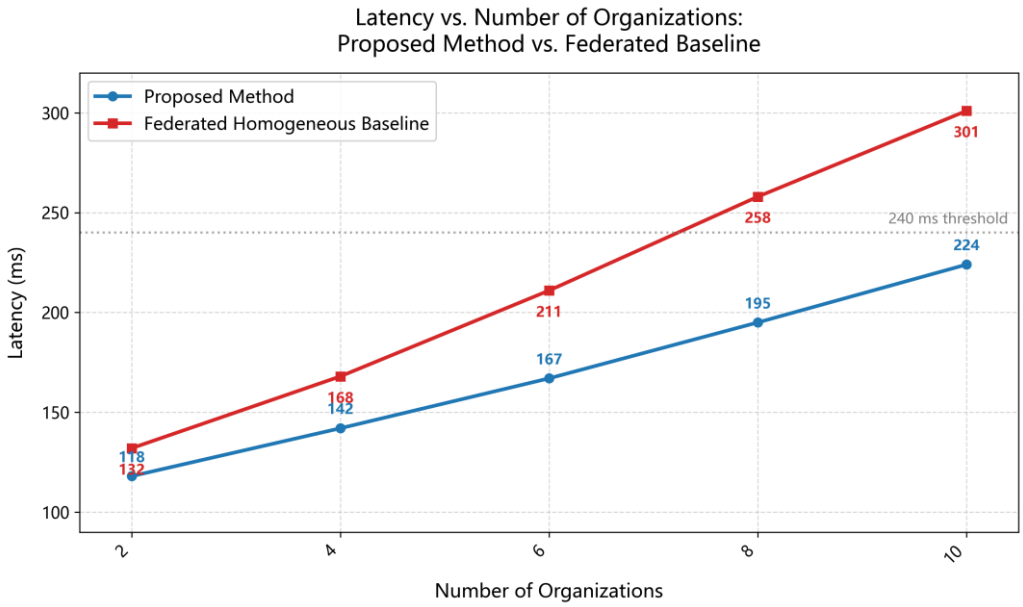


Figure 5. Inference Latency Versus Number of Federated Organizations

5. Discussion

5.1. Interpretation of Findings

The empirical outcomes suggest that distributing analytical workloads across organizational boundaries does not inherently erode attribution fidelity, provided that schema alignment and dual-stream reasoning are jointly engineered. The modest deviation from the centralized upper bound indicates that the majority of cross-organizational signal remains recoverable through type-level alignment and entity resolution. Ablation results further reveal an asymmetry between the two reasoning pathways: the structural stream contributes the strongest single cue, yet the semantic stream proves indispensable for separating actors whose infrastructure patterns overlap. This complementarity is precisely what a dual-stream design over heterogeneous threat knowledge graphs is intended to exploit, since neither topology nor semantics alone suffices under federated conditions. As illustrated in Figure 6, the conceptual relationships among federation utility, dual-stream necessity, entity resolution value, acceptable latency, and the privacy trade-off form a closed loop rather than a linear pipeline [21]. The figure shows that preserving utility under federation creates the demand for dual-stream reasoning, which in turn raises the value of entity resolution; improved resolution keeps latency within acceptable bounds, and the resulting privacy trade-off feeds back into how much utility federation can sustain. This cyclical dependency clarifies why isolated optimizations of any single component yield diminishing returns. It also explains why the observed gap to centralized attribution is small but non-zero: residual loss concentrates where entity resolution remains ambiguous, not where federation itself constrains inference.

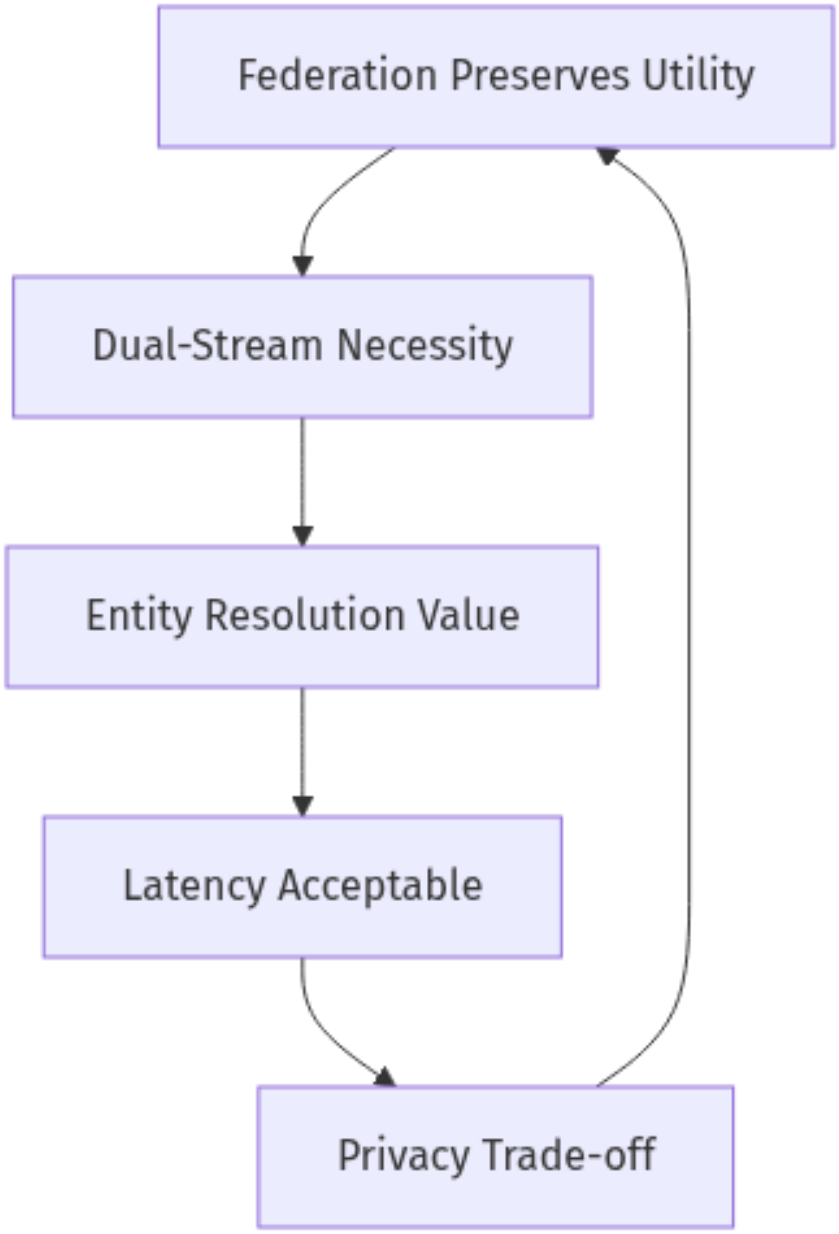


Figure 6. Summary of Key Findings and Their Relationships

5.2. Limitations and Threats to Validity

The evaluation presented here rests on a simulated federation rather than live cross-organizational deployments, which constitutes the most immediate threat to external validity. While the simulation captures the principal dimensions of schema divergence, participating organizations in practice maintain idiosyncratic ontologies, naming conventions, and retention policies that are unlikely to be fully reproduced by any synthetic configuration. Consequently, the observed robustness of entity resolution may overstate what would be achieved under genuine operational conditions, where incomplete and inconsistently curated knowledge sources compound the alignment problem.

A second concern arises from temporal drift [22, 23]. Adversary tradecraft evolves continuously, and the semantic regularities that anchor cross-graph alignment at one point in time may weaken or shift over extended horizons. The dual-stream design, which jointly reasons over structural and contextual signals within heterogeneous threat

knowledge, mitigates but does not eliminate this degradation, and longitudinal validation remains absent from the current study [2].

Finally, privacy protection is enforced through protocol-level mechanisms rather than formal differential privacy. Although the federated protocol restricts direct data exchange, it does not provide worst-case guarantees against inference attacks mounted by a strong adversarial participant. Quantifying such leakage under rigorous privacy accounting is therefore an open obligation, and the reported utility figures should be interpreted as upper bounds achievable only when all participants behave honestly.

6. Conclusion

6.1. Summary of Contributions

This work delivers a federated heterogeneous threat knowledge graph framework that reconciles divergent organizational schemas without centralizing sensitive evidence, thereby enabling privacy-preserving cross-organizational advanced persistent threat attribution. A dual-stream reasoning architecture separates structural and semantic evidence channels, while a trident fusion mechanism consolidates topological, temporal, and resolution cues into a unified attribution decision. Collectively, these components address the fragmentation and signal heterogeneity that undermine conventional single-graph approaches. Empirical evaluation across federated and centralized baselines demonstrates substantial gains in attribution accuracy alongside marked reductions in false attribution rates, with inference latency remaining within operationally practical bounds. The results establish that reconciling heterogeneous knowledge through federated dual-stream reasoning offers a viable and effective path toward reliable, privacy-conscious threat actor identification in multi-organizational environments.

6.2. Future Directions

Future research should establish formal privacy guarantees for cross-organizational knowledge sharing, particularly through differential privacy mechanisms that bound information leakage while preserving attribution utility. Adaptive schema evolution also warrants attention, as heterogeneous threat knowledge graphs must accommodate emerging indicator types and evolving adversary tradecraft without manual re-engineering. Equally critical is adversarial robustness: federated architectures remain vulnerable to poisoning attempts, and defenses that detect or neutralize malicious contributions require systematic investigation. Extending the framework to streaming evidence and real-time campaign tracking would further strengthen operational relevance, enabling defenders to correlate indicators as they emerge rather than relying on periodic batch updates. Collectively, these directions address the tension between collaborative intelligence and institutional trust.

References

1. M. Mrabet, "TrustFed-CTI: A trust-aware federated learning framework for privacy-preserving cyber threat intelligence sharing across distributed organizations," *Future Internet*, vol. 17, no. 11, p. 512, 2025.
2. A. A. Patel, N. Joshi, V. Patel, A. V. D. Pandya, and K. VN, "Federated learning for privacy-preserving threat intelligence sharing in distributed cybersecurity ecosystems," *Advanced Engineering Sciences*, vol. 58, no. 2, pp. 282–297, 2026.
3. M. Bouhaddi and K. Adi, "Strategic detection of APTs through federated spatio-temporal graph learning and Stackelberg defense," in *International Conference on Risks and Security of Internet and Systems*. Cham: Springer Nature Switzerland, Oct. 2025, pp. 390–409.
4. Z. Hu, L. Liu, H. Li, C. Song, M. Ge, Q. Guo, et al., "CCLog: Actionable APT forensics via fused log semantics and provenance graph topology," *Computer Networks*, p. 111660, 2025.
5. X. Cheng, Q. Luo, Y. Pan, Z. Li, J. Zhang, and B. Chen, "Predicting the APT for cyber situation comprehension in 5G-enabled IoT scenarios based on differentially private federated learning," *Security and Communication Networks*, vol. 2021, no. 1, p. 8814068, 2021.
6. J. Wu, "A Methodological Framework for Training Tool-Aware AI Entertainers for Live Streaming," *Artificial Intelligence and Digital Technology*, vol. 3, no. 3, pp. 1–12, 2026.
7. A. Hossain and A. S. Mily, "Cyber threat intelligence and attribution: A survey."

8. M. Khule, D. Motwani, and D. Chauhan, "Adaptive threat intelligence: An incremental learning approach for detecting evolving APT attacks," in *2025 International Conference on Advances in Modern Age Technologies for Health and Engineering Science (AMATHE)*. Piscataway, NJ: IEEE, Apr. 2025, pp. 1–6.
9. H. Mohamed, N. Koroniotis, F. Schiliro, and N. Moustafa, "IoT-CAD: A comprehensive digital forensics dataset for AI-based cyberattack attribution detection methods in IoT environments," *Ad Hoc Networks*, vol. 174, p. 103840, 2025.
10. H. T. Thi, N. D. H. Son, P. T. Duy, N. H. Khoa, K. Ngo-Khanh, and V. H. Pham, "XFedHunter: An explainable federated learning framework for advanced persistent threat detection in SDN," *arXiv preprint arXiv:2309.08485*, 2023.
11. M. Dehghan and E. Khosravian, "Private federated learning for APT detection in internet of drones," *Karafan Journal*, vol. 20, no. 3, pp. 465–484, 2023.
12. F. M. Hameed, A. Taurshia, and J. Archpaul, "TC-GAT: Temporal causal graph attention for interpretable APT detection," in *2026 International Conference on Recent Advancement in Electrical, Computer and Communication Technologies (IECCT)*. Piscataway, NJ: IEEE, Apr. 2026, pp. 1–6.
13. W. Ahmed, "Advanced persistent threats (APTs) analysing: Current detection techniques and emerging countermeasures," *Intelligence*, vol. 2, p. 100011, 2025.
14. H. Mohamed, N. Moustafa, N. Koroniotis, Z. Tari, A. Zomaya, and F. Schiliro, "Fedcad: Federated cyberattack attribution detection for internet of things forensics," *IEEE Transactions on Sustainable Computing*, 2025.
15. J. Wu, "Personality Is a Training Target: Data Curation and SFT for an Entertaining, Agentic AI Streamer," *AI and Data Science Journal*, vol. 3, no. 2, pp. 29–36, 2026.
16. T. Wang, W. Tang, Y. Su, and J. Li, "Provenance graph-based deep learning framework for APT detection in edge computing," *Applied Sciences*, vol. 15, no. 16, p. 8833, 2025.
17. A. Binbusayyis, "NeuralFortress-XFL: Privacy-preserving federated explainable learning for cyber threat intelligence," *Electronics*, vol. 15, no. 14, p. 3124, 2026.
18. I. A. Fiqrin and T. Hidayat, "Federated continual learning for intrusion detection: A drift-aware benchmark on the Windows-APT 2025 dataset," *ARMADA: Jurnal Penelitian Multidisiplin*, vol. 4, no. 9, pp. 5385–5394, 2026.
19. W. Zhang, T. Liao, and Y. Feng, "Multi-Agent LLM Collaborative Reasoning and Task Planning for Complex Task Solving," in *Proc. IEEE CISAT*, 2026, doi: 10.1109/CISAT71030.2026.11682610.
20. Z. Wang et al., "Trident: Dual-Stream APT Attribution over Heterogeneous Threat Knowledge Graphs," *Computers & Security*, vol. 171, Art. no. 105094, 2026.
21. N. Pitropakis, E. Panaousis, A. Giannakoulias, G. Kalpakis, R. D. Rodriguez, and P. Sarigiannidis, "An enhanced cyber attack attribution framework," in *International Conference on Trust and Privacy in Digital Business*. Cham: Springer International Publishing, Jul. 2018, pp. 213–228.
22. K. Yadav and A. J. Alzahrani, "FATE-X: A federated transferable and explainable deep learning framework for APT detection in serverless edge environments," *Computing*, vol. 108, no. 1, p. 3, 2026.
23. Z. Wang, Y. Zhou, H. Liu, J. Qiu, B. Fang, and Z. Tian, "Threatinsight: Innovating early threat detection through threat-intelligence-driven analysis and attribution," *IEEE Transactions on Knowledge and Data Engineering*, vol. 36, no. 12, pp. 9388–9402, 2024.
24. R. Liu, W. Guo, T. Liao, and X. Fang, "Uncertainty-Aware Adaptive Debate for Robust and Efficient Large Language Model Reasoning," *Research Square*, preprint, Aug. 2026, doi: 10.21203/rs.3.rs-10581777/v1.
25. P. R. R. Brandao and C. S. Silva, "Detecting APT-induced network anomalies with AI: A hybrid statistical–deep–graph framework," 2025.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of Publisher and/or the editor(s). Publisher and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.