SecGuard: Multi-Dimensional Provenance Analysis for Self-Supervised APT Detection

Chen CHEN, Yunchun LI, Mingyuan XIA, Wei LI

Frontiers of Computer Science, DOI: 10.1007/s11704-025-50797-z

Problems & Ideas

- Problems of conventional provenance-based approaches:
  - Vulnerable to mimicry attacks where malicious actions are disguised as benign.
  - Lack fine-grained detection and overlook critical semantic and temporal signals in system audit logs.
- Ideas: A self-supervised framework (SecGuard) that integrates structural, semantic, and temporal features for robust, fine-grained APT detection.

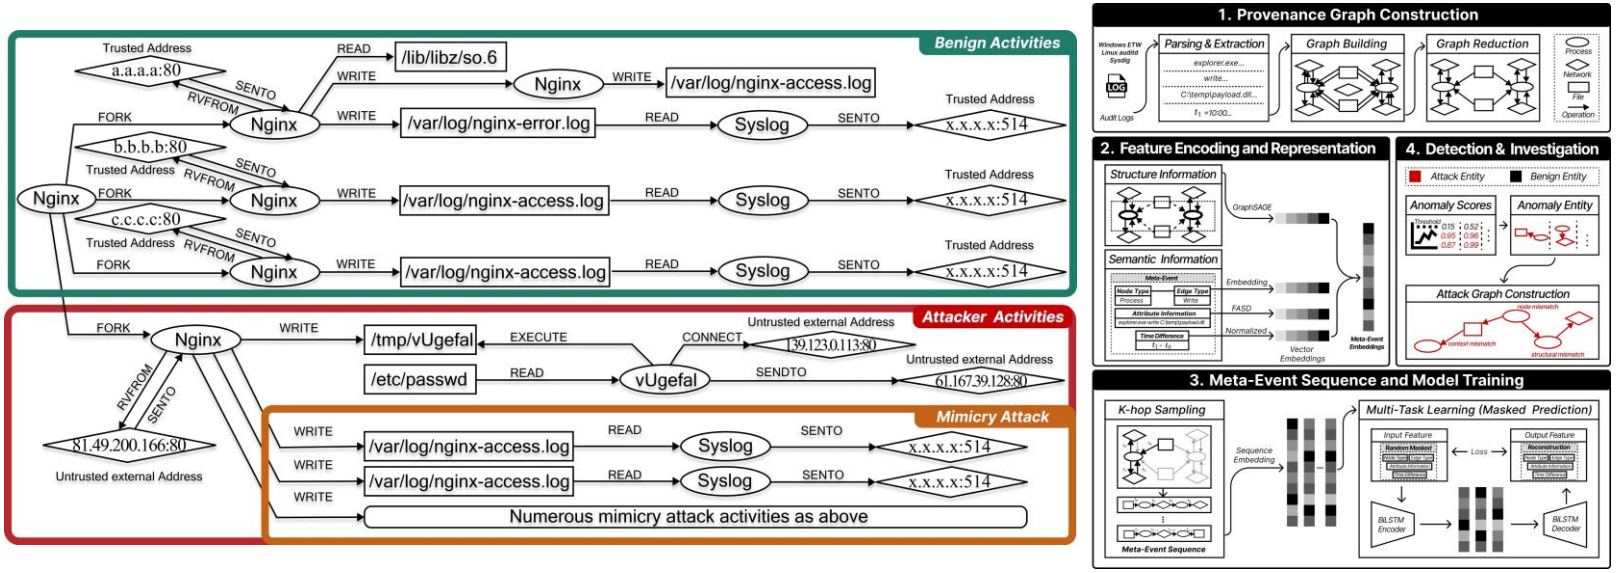

An APT scenario showing malicious (red) and mimicry (orange) activities interwoven with benign behaviors (green), illustrating the detection challenge.

Main Contributions

- Contributions:
  - A novel self-supervised APT detection framework that uniquely integrates structural, semantic, and temporal analysis of provenance graphs.
  - A multi-dimensional approach that provides strong robustness against sophisticated mimicry attacks by making evasion significantly more difficult.
  - Ablation studies confirm that each core component, including our FASD algorithm, is critical for achieving efficient and accurate detection.

Table 1 Performance comparison of SecGuard, ThreaTrace, and FLASH on entity level APT detection across DARPA TC datasets and SimAttack datasets. Pre.: Precision; Rec.: Recall; F1.:F1-score.

Datasets	SecGuard				ThreaTrace [2]				FLASH [4]			
	TP/TN/FP/FN	Pre.	Rec.	F1.	TP/TN/FP/FN	Pre.	Rec.	F1.	TP/TN/FP/FN	Pre.	Rec.	F1.
THEIA (E3)	25 323/3 503 938/1399/28	**0.95**	0.99	**0.97**	25 297/3 501 561/3765/65	0.87	0.99	0.93	25 318/3 503 451/1875/44	0.93	0.99	0.96
THEIA (E5)	150 816/8 345 915/38 459/12 019	**0.80**	0.92	**0.85**	150 286/8 321 358/63 137/12 428	0.70	0.92	0.80	150 713/8 333 932/50 239/12 325	0.75	0.92	0.83
Trace (E3)	67 361/2 413 301/2727/1	**0.96**	0.99	**0.98**	67 382/2 389 233/26 774/1	0.72	0.99	0.83	67 382/2 412 202/3805/1	0.95	0.99	0.97
Trace (E5)	25 091/1 823 599/2453/3520	**0.91**	**0.87**	**0.89**	25 024/1 819 337/6014/4288	0.81	0.85	0.83	25 053/1 820 419/5314/3877	0.83	0.86	0.84
CADETS (E3)	12 812/706 644/358/4	**0.97**	0.99	**0.98**	12 848/705 605/1361/4	0.90	0.99	0.95	12 851/706 246/720/1	0.95	0.99	0.97
CADETS (E5)	17 733/478 913/4209/1995	**0.81**	**0.89**	**0.85**	17 685/472 045/10 281/2839	0.63	0.86	0.73	17 694/476 963/5598/2595	0.76	0.87	0.81
SimAttck-1	14 353/695 789/1345/5	**0.91**	0.99	**0.95**	14 224/695 021/2229/18	0.86	0.99	0.92	14 249/695 521/1715/7	0.89	0.99	0.94
SimAttck-2	13 183/803 678/1475/27	**0.90**	0.99	**0.94**	13 021/802 931/2367/44	0.85	0.99	0.91	13 143/803 437/1752/31	0.88	0.99	0.93

Table 2 Performance comparison of SecGuard, ThreaTrace, Flash, and Unicorn on subgraph-level APT detection across Unicorn dataset scenarios (SC-1, SC-2, and Wget Pre.: Precision; Rec.: Recall; F1.:F1-score.

Datasets	System	Pre.	Rec.	F1.
Unicorn SC-1 [1]	Unicorn [1]	0.85	0.96	0.90
	Flash [4]	0.92	0.96	0.94
	ThreaTrace [2]	0.93	0.98	0.95
	SecGuard	**0.95**	0.98	**0.97**
Unicorn SC-2 [1]	Unicorn	0.75	0.80	0.78
	Flash	0.96	0.96	0.96
	ThreaTrace	0.91	0.96	0.93
	SecGuard	**0.93**	**0.98**	0.96
Unicorn Wget [1]	Unicorn	0.86	0.95	0.90
	Flash	0.97	0.93	0.96
	ThreaTrace	0.93	0.98	0.95
	SecGuard	**0.97**	**0.98**	**0.98**

Experimental results show SecGuard consistently outperforms baseline methods on DARPA and Unicorn datasets in both entity-level and subgraph-level detection.