

# KnowHow: Automatically Applying High-Level CTI Knowledge for Interpretable and Accurate Provenance Analysis

Yuhan Meng<sup>†</sup>, **Shaofei Li<sup>†</sup>**, Jiaping Gui<sup>‡</sup>, Peng Jiang<sup>§</sup>, and Ding Li<sup>†\*</sup>

<sup>†</sup> Key Laboratory of High-Confidence Software Technologies (MOE), School of Computer Science, Peking University,

<sup>‡</sup> School of Computer Science, Shanghai Jiao Tong University, <sup>§</sup> Southeast University



北京大學  
PEKING UNIVERSITY



上海交通大學  
SHANGHAI JIAO TONG UNIVERSITY



東南大學  
SOUTHEAST UNIVERSITY

# Advanced Persistent Threats

---

- **APT attacks** pose a critical challenge to the modern world.
  - **Advanced:** Use cutting-edge techniques → Zero-day exploit, well-designed malware.
  - **Stealthy:** Evade detection by operating covertly → Living-off-the-land (LotL) attack, fileless attack.
  - **Persistent:** Maintain long-term access. → The SolarWinds attackers remained undetected for over 14 months.
- This underscores the critical need for effective detection.



# Provenance-Based Detection Systems

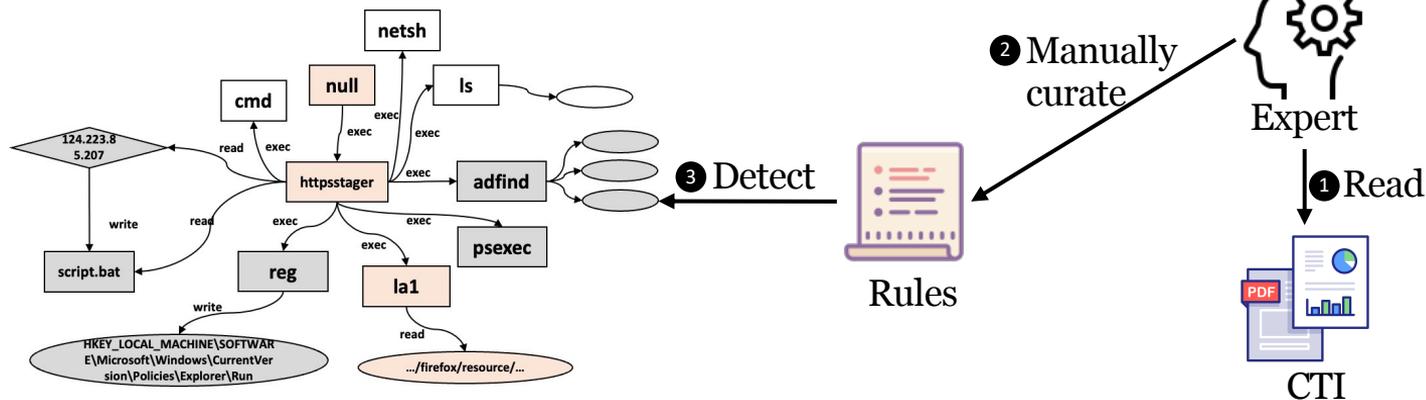
---



- Provenance-based detection systems build a directed, attributed provenance graph that models system behaviors over time.
  - **Node:** system entities (processes, files, sockets, etc.)
  - **Edge:** system events (read, write, fork, execve, sendto, recvfrom, etc.)

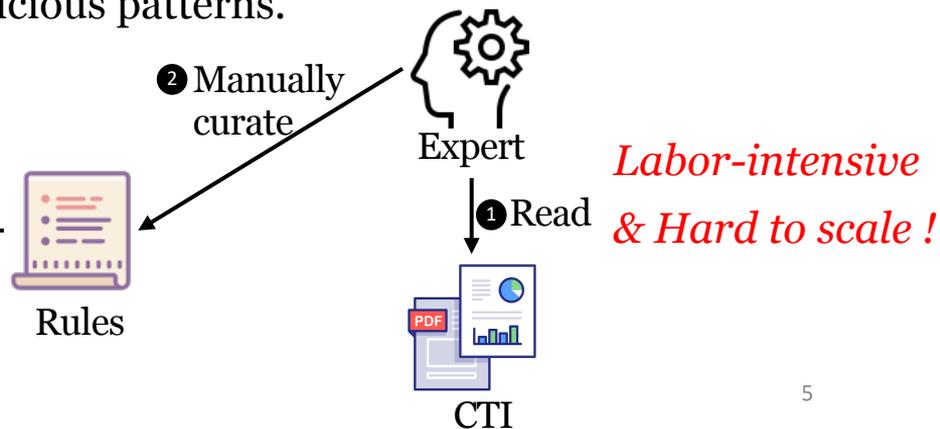
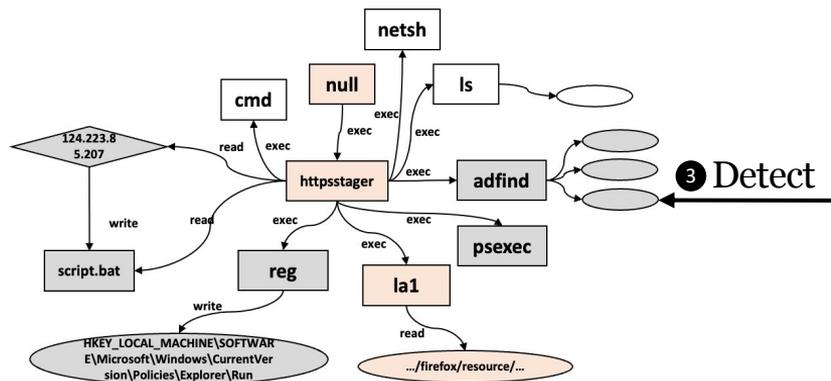
# Provenance-Based Detection Systems

- Provenance-based detection systems build a directed, attributed provenance graph that models system behaviors over time.
  - Node:** system entities (processes, files, sockets, etc.)
  - Edge:** system events (read, write, fork, execve, sendto, recvfrom, etc.)
- One effective class of methods is **knowledge-driven approaches**
  - Encode attack knowledge, **manually** curated by security experts from **Cyber Threat Intelligence (CTI)**, to detect malicious patterns.



# Provenance-Based Detection Systems

- Provenance-based detection systems build a directed, attributed provenance graph that models system behaviors over time.
  - Node:** system entities (processes, files, sockets, etc.)
  - Edge:** system events (read, write, fork, execve, sendto, recvfrom, etc.)
- One effective class of methods is **knowledge-driven approaches**
  - Encode attack knowledge, manually curated by security experts from Cyber Threat Intelligence (CTI), to detect malicious patterns.



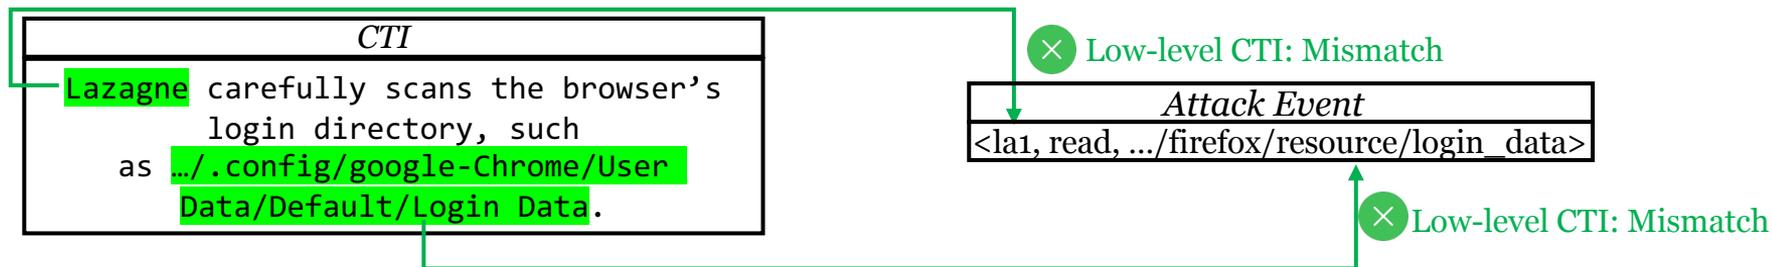
# Motivation

---

- What prevents knowledge-driven approaches from being automatic and extensible?

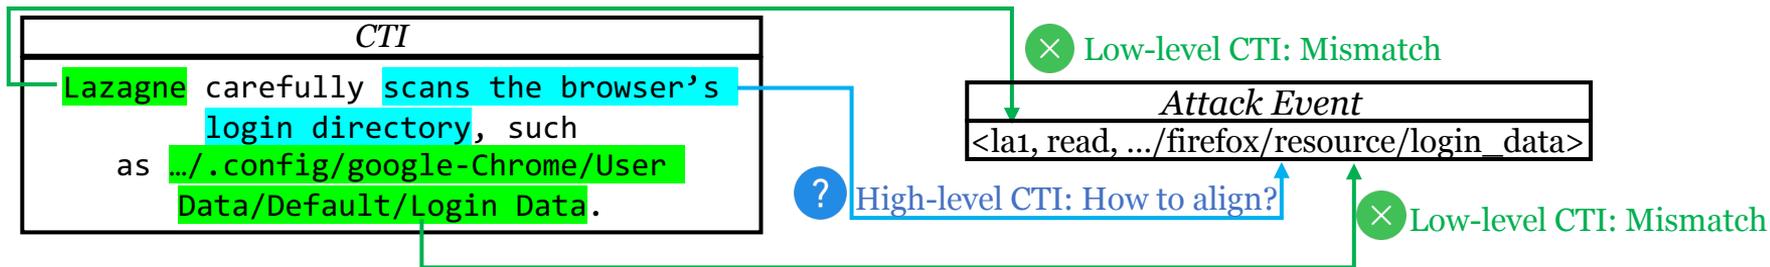
# Motivation

- What prevents knowledge-driven approaches from being automatic and extensible?
- Two types of CTI knowledge:
  - **Low-level** (instantiated) : Indicator of Compromise (IoC) in CTI → Automatically extracted and matched by IoC-based methods (e.g., EXTRACTOR, TTPDrill, AttacKG, LADDER).



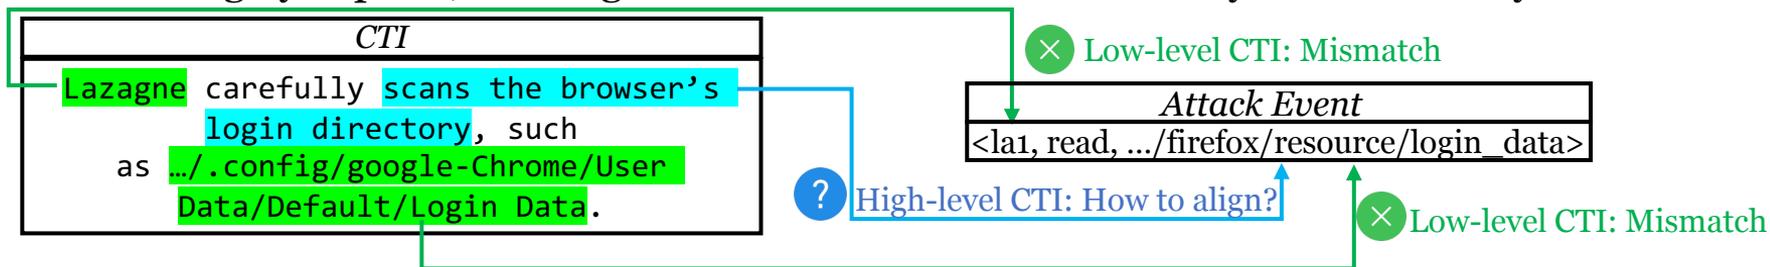
# Motivation

- What prevents knowledge-driven approaches from being automatic and extensible?
- Two types of CTI knowledge:
  - **Low-level** (instantiated) : Indicator of Compromise (IoC) in CTI → Automatically extracted and matched by IoC-based methods (e.g., EXTRACTOR, TTPDrill, AttacKG, LADDER).
  - **High-level** (non-instantiated) : General behavioral descriptions in CTI → Cannot be matched automatically due to the **lack of linguistic similarity** with provenance events. (e.g., LADDER fails to correlate CTI sentences with system events).



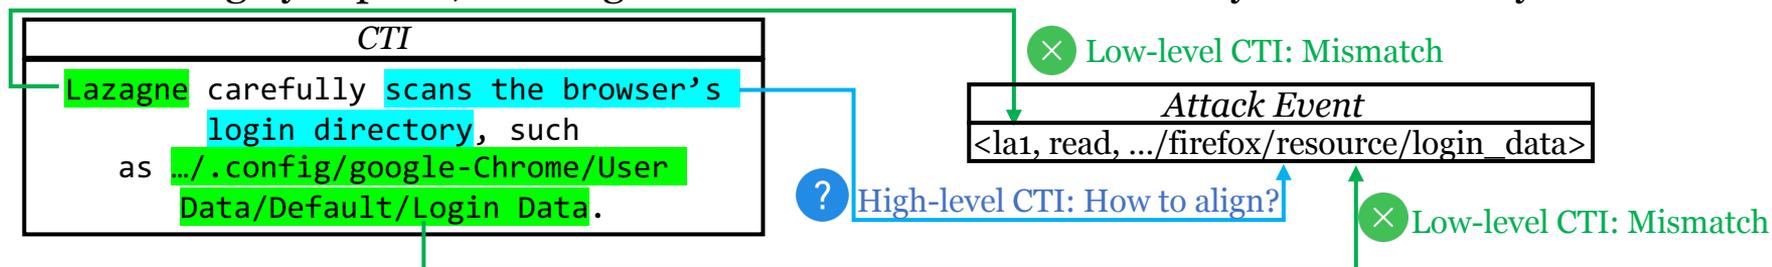
# Motivation

- What prevents knowledge-driven approaches from being automatic and extensible?
- Two types of CTI knowledge:
  - **Low-level** (instantiated) : Indicator of Compromise (IoC) in CTI → Automatically extracted and matched by IoC-based methods (e.g., EXTRACTOR, TTPDrill, AttacKG, LADDER).
  - **High-level** (non-instantiated) : General behavioral descriptions in CTI → Cannot be matched automatically due to the **lack of linguistic similarity** with provenance events. (e.g., LADDER fails to correlate CTI sentences with system events).
- Consequently, this rich high-level knowledge can **only be leveraged through manual rule** writing by experts, limiting the automation and extensibility of detection systems.



# Motivation

- What prevents knowledge-driven approaches from being automatic and extensible?
- Two types of CTI knowledge:
  - **Low-level** (instantiated) : Indicator of Compromise (IoC) in CTI → Automatically extracted and matched by IoC-based methods (e.g., EXTRACTOR, TTPDrill, AttacKG, LADDER).
  - **High-level** (non-instantiated) : General behavioral descriptions in CTI → Cannot be matched automatically due to the **lack of linguistic similarity** with provenance events. (e.g., LADDER fails to correlate CTI sentences with system events).
- Consequently, this rich high-level knowledge can **only be leveraged through manual rule** writing by experts, limiting the automation and extensibility of detection systems.



- **Observation: The Semantic Gap Prevents Automatic Use of High-Level CTI in Knowledge-based Approaches**

# Key Insight

---

- To overcome this semantic gap, our key insight is that high-level CTI can be automatically mapped to provenance events by **aligning them into a subject-verb-object (SVO) semantic structure**.

# Key Insight

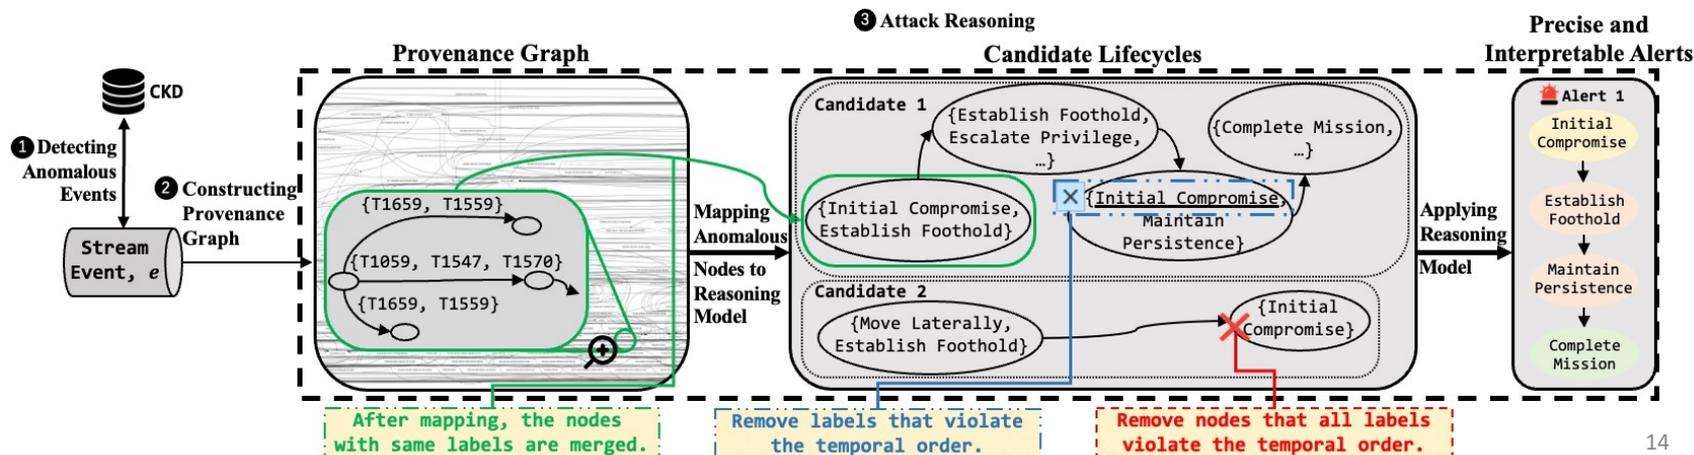
- To overcome this semantic gap, our key insight is that high-level CTI can be automatically mapped to provenance events by **aligning them into a subject-verb-object (SVO) semantic structure**.
- CTI descriptions can be extracted into three semantic roles, which can map to the key elements in provenance events:
  - Attack conductors (who) → processes
  - Attack actions (what operation) → syscall
  - Attack targets (on what) → targeted files
- Unlike static and context-free IoCs, this structure captures both entities and their contextual descriptions, enabling more accurate, context-aware detection.

Sentences in CTI	Provenance Events
<u>Lazagne</u> <u>scanned</u> <u>browser resources</u> to steal credentials.	{<left out> "evt.type": "read", "fd.name": " <u>_/firefox/resource...</u> ", "proc.cmdline": "la1 -s firefox", "proc.name": " <u>la1</u> ", <left out>}
Attackers may <u>send</u> the <u>stolen data</u> to <u>the external server</u> hold held by them, to achieve the goal of exfiltration.	{<left out> "evt.type": "sendto", "fd.name": " <u>17.x.x.x-&gt;12.5.8.10</u> ", "proc.cmdline": "scp <u>./collect</u> <u>/12.5.8.10</u> ", "proc.name": "scp", <left out>}
Some <u>malicious js scripts</u> may be executed and <u>injected</u> into <u>proc filesystem</u> , which can evade process-based defenses as well as possibly elevate privileges.	{<left out> "evt.type": "write", "fd.name": " <u>/proc/filesystem</u> ", "proc.cmdline": " sh -r <u>./exp.js</u> ", "proc.name": " sh", <left out>}

- **C1: How to extract complete and context-preserving SVO triples from CTI reports.**
  - **Problem:** Real-world CTI uses complex, passive, or fragmented sentences, while standard NLP parsers may miss critical modifiers of the entities.
  - **Solution1:** Propose general Indicator of Compromise (gIoC) representation with context-aware SVO extraction.
- **C2: How to efficiently map gIoCs to system events.**
  - **Problem:** Terms like “steal” or “exfiltrate” in gIoCs have no lexical overlap with system entities like “read” or “sendto”, making direct matching impossible.
  - **Solution2:** Design CTI Knowledge Database (CKD) to enable semantic lifting and accelerated mapping against system events.
- **C3: How to detect multi-stage APT behaviors from isolated event-level mapping.**
  - **Problem:** Event-level detection yields fragmented alerts, while APTs require reconstructing cross-process, time-spanning attack chains.
  - **Solution3:** Novel attack reasoning method based on the stages of the APT Lifecycle.

# Our Solution: KnowHow

- KnowHow bridges the semantic gap between high-level CTI and low-level provenance events and achieves attack detection through a three-stage pipeline.
  - Preprocessing: Extract gIoCs from CTI reports, and store them in the CKD.
  - Stage 1: Detect anomalous system events by querying gIoC in CKD.
  - Stage 2: Construct the provenance graph starting from the detected anomalous events.
  - Stage 3: Apply reasoning model to analyze candidate alert lifecycles, to eliminate false positives.



# C1: How to Extract High-level Knowledge

---

- **Solution1:** We propose general Indicator of Compromise (gIoC) representation with context-aware SVO extraction.
- **Definition of gIoC**
  - We define gIoC as a structured, semantic triple that captures the core behavioral intent of an attack described in CTI, in the form of:
$$gIoC = \langle Conductor, Action, Target \rangle$$
  - The key difference between gIoCs and conventional SVO triples is that the subjects and objects of gIoCs are **attack-relevant concepts** and their **associated information**, such as modifiers describing these concepts.
  - gIoCs extract attack information from **a behavioral perspective**, instead of the instance-level details, enabling it to reflect higher-level information.

# C1: How to Extract High-level Knowledge

---

- **Solution1:** We propose gIoC (general Indicator of Compromise) representation with context-aware SVO extraction.
- **Extraction of gIoC**
  - Step 1: Identify **attack-relevant concepts ( $N$ )**, if  $N$  fulfills one of the following:
    - $N$  is an IoC, such as a file name, an IP address, a file hash, etc.
    - $N$  is a domain name.
    - $N$  is the name of an application or malware.
    - $N$  is a command (e.g., *cp*) or its full name (e.g., *copy*)
    - $N$  represents a general concept of system objects, including but not limited to terms like “file,” “directory,” “IP address,” “process,” “application,” “registry,” and their synonyms.

# C1: How to Extract High-level Knowledge

---

- **Solution1:** We propose gIoC (general Indicator of Compromise) representation with context-aware SVO extraction.
- **Extraction of gIoC**
  - Step 1: Identify **attack-relevant concepts ( $N$ )**, if  $N$  fulfills one of the following:
    - $N$  is an IoC, such as a file name, an IP address, a file hash, etc.
    - $N$  is a domain name.
    - $N$  is the name of an application or malware.
    - $N$  is a command (e.g., *cp*) or its full name (e.g., *copy*)
    - $N$  represents a general concept of system objects, including but not limited to terms like “file,” “directory,” “IP address,” “process,” “application,” “registry,” and their synonyms.
  - Step 2: Identify the **associated information** of the attack-relevant concepts
    - We extract the **modifiers** (e.g., “browser data”) and **subclauses** (e.g., “steal the credential files of system users”) that describe the concept and serve as its associated information.

## C2: How to Efficiently Use gIoCs

---

- **Solution2:** We design CTI Knowledge Database (CKD) to enable semantic lifting and accelerated mapping when using gIoCs to detect system events.
- **Structure of CKD**
  - CKD is defined as a set of **ATT&CK Technique Information Entries** (ATIEs), each representing an attack technique as defined by MITRE ATT&CK.
  - An ATIE comprises four fields:
    - a unique ID for a technique in ATT&CK (*uid*),
    - a description of the technique (*des*),
    - a CTI list (*list<sub>cti</sub>*),
    - and a **gIoC list** (*list<sub>gioc</sub>*).
  - *list<sub>gioc</sub>* contains gIoCs, which capture high-level attack knowledge reported in the CTI reports in *list<sub>cti</sub>*.

## C2: How to Efficiently Use gIoCs

---

- **Solution2:** We design CTI Knowledge Database (CKD) to enable semantic lifting and accelerated mapping when using gIoCs to *detect system events*.
- **Operation of CKD: *ProvQ***
  - *ProvQ* matches the low-level system events to ATIEs in CKD, which are then flagged as **anomalous**.
  - For input system event  $e = (source, destination, syscalltype, commandline)$ , *ProvQ* returns a list of ATIEs that **match** the given event.
  - An event  $e$  matches an ATIE  $t$  only if its similarity score  $Sim(e, t) > \theta_q$ , where  $\theta_q$  is the given query threshold. And  $Sim(e, t)$  is defined as

$$Sim(e, t) = \sum S(e.y, t), \text{ where } y \in \{source, destination, syscalltype, commandline\},$$

and  $S(e.y, t)$  is the occurrences of gIoCs in  $t$  that *appear in*  $e.y$  field.

## C2: How to Efficiently Use gIoCs

---

- **Solution2:** We design CTI Knowledge Database (CKD) to enable semantic lifting and accelerated mapping when using gIoCs to *detect system events*.
- **Operation of CKD: *ProvQ***
  - *ProvQ* matches the low-level system events to ATIEs in CKD, which further is flagged as **anomalous**.
  - For input system event  $e = (source, destination, syscalltype, commandline)$ , *ProvQ* returns a list of ATIEs that **match** the given event.
  - An event  $e$  matches an ATIE  $t$  only if its similarity score  $Sim(e, t) > \theta_q$ , where  $\theta_q$  is the given query threshold. And  $Sim(e, t)$  is defined as

$$Sim(e, t) = \sum S(e.y, t), \text{ where } y \in \{source, destination, syscalltype, commandline\},$$

and  $S(e.y, t)$  is the occurrences of gIoCs in  $t$  that *appear in*  $e.y$  field.

*How to define 'appear in'?*

## C2: How to Efficiently Use gIoCs

---

- **Solution2:** We design CTI Knowledge Database (CKD) to enable *semantic lifting* and accelerated mapping when using gIoCs to detect system events.
- **Operation of CKD: *ProvQ***
  - Calculate  $S(e, y, t)$  : Terms like “steal” or “exfiltrate” in gIoC in  $t$  have **no lexical overlap** with system entities like “read” or “sendto” in  $e, y$ , making direct matching impossible.

# C2: How to Efficiently Use gIoCs

• **Solution2:** We design CTI Knowledge Database (CKD) to enable *semantic lifting* and accelerated mapping when using gIoCs to detect system events.

## • Operation of CKD: *ProvQ*

- Calculate  $S(e, y, t)$  : Terms like “steal” or “exfiltrate” in gIoC in  $t$  have **no lexical overlap** with system entities like “read” or “sendto” in  $e, y$ , making direct matching impossible.
- We propose **semantic lifting** to elevate low-level system events to natural language sentences that can be matched with gIoCs, which further helps to calculate  $S(e, y, t)$  precisely.

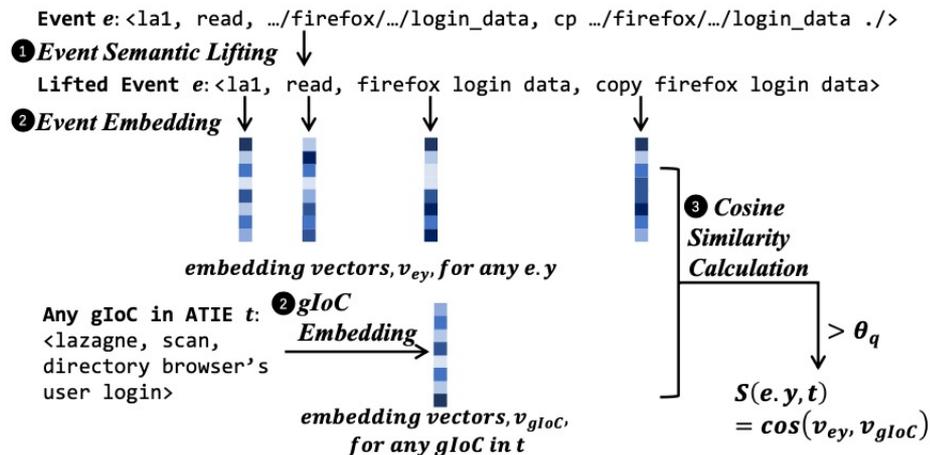
	System Identifier	Lifted Sentences
Linux File	/etc/D/*/F.E	etc D E file
	/var/D/*/F.E	var D E file
	/proc/[PID]/D/*/F.E	proc D E file
	/bin/D/*/F.E, /sbin/D/*/F.E, /usr/bin/D/*/F.E, /usr/sbin/D/*/F.E, /usr/local/bin/D/*/F.E, /usr/local/sbin/D/*/F.E	F E file
	/home/aa/D/ * /F.E\$	user D F E file
	/root/D/*/F.E	root user D F E file
	/lib/D/*/F.E, /lib32/D/*/F.E, /lib64/D/*/F.E, /usr/local/lib/D/*/F.E, /xx/lib/D/*/F.E	D library file
	other: */F.E	E file
	Windows File	HKEY *, HKCU*, HKCR \*, HKLM*, HKU*, HKCC*
c:\\windows\system32\D\*F.E		windows system D F.E file
c:\\windows\D\*F.E		windows system D F.E file
c:\\ProgramFiles\D\*F.E, c:\\ProgramFiles(x86)\\D\*F.E		D F E file
other: */F.E		F E file

## C2: How to Efficiently Use gIoCs

- **Solution2:** We design CTI Knowledge Database (CKD) to enable *semantic lifting* and accelerated mapping when using gIoCs to detect system events.

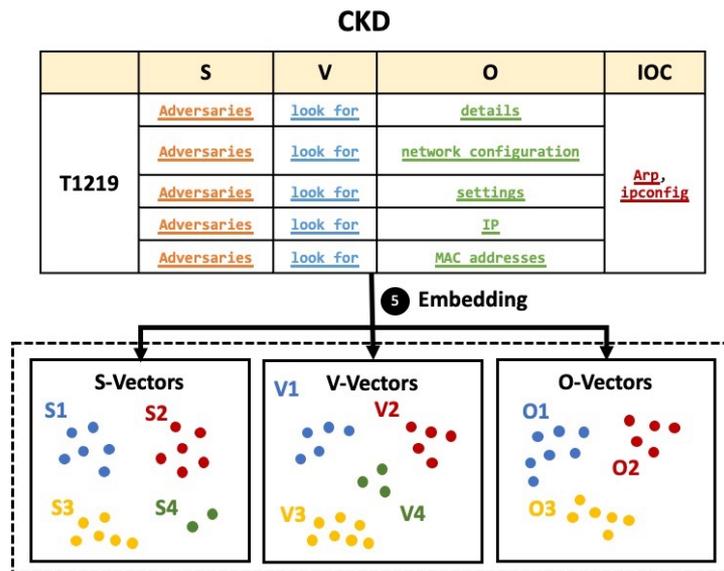
- **Operation of CKD: *ProvQ***

- Calculate  $S(e, y, t)$ : Terms like “steal” or “exfiltrate” in gIoC in  $t$  have **no lexical overlap** with system entities like “read” or “sendto” in  $e, y$ , making direct matching impossible.
- We propose **semantic lifting** to elevate low-level system events to natural language sentences that can be matched with gIoCs, which further helps to calculate  $S(e, y, t)$  precisely.



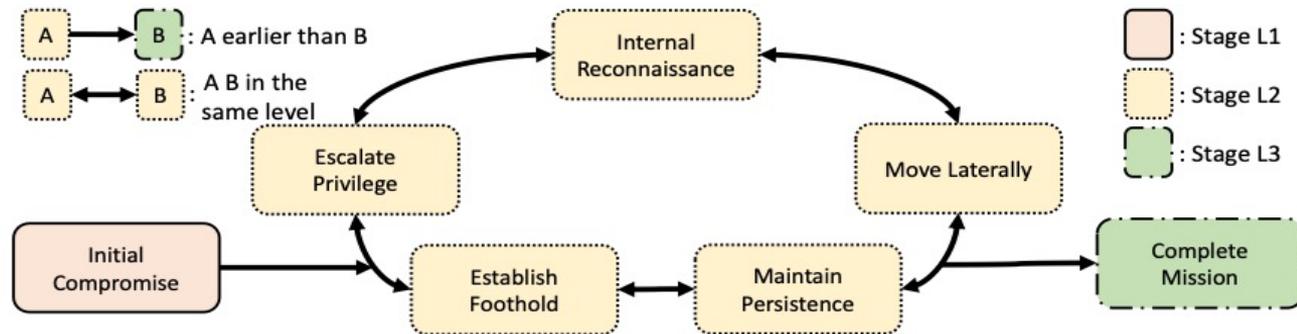
## C2: How to Efficiently Use gIoCs

- **Solution2:** We design CTI Knowledge Database (CKD) to enable semantic lifting and *accelerated* mapping when using gIoCs to detect system events.
- **Accelerating *ProvQ*:**
  - We propose a **two-stage search method** to accelerate the speed of *ProvQ*
  - **Offline stage:** Cluster all gIoCs using Mean-Shift algorithm on their embedding vectors
  - **Online stage:** For event field  $e, y$ ,
    - we first find the closest gIoC cluster
    - then compute similarity only within that cluster
  - **Result:** Avoids over 70% unnecessary comparisons while preserving recall.



# C3: How to Detect Multi-stage APTs

- **Solution3:** Novel attack reasoning method based on the stages of the APT Lifecycle
- **Reasoning Model: A relaxed APT Lifecycle model**
  - **Lifecycle Completeness:** The attack must have initial compromise step and one of the steps in stage L2.
  - **Lifecycle Temporal Order:** The attack step in later stage must happen later than the step in previous stage.
  - We design a **one-to-limited mapping** optimization to map anomalous nodes to reasoning model.



## • Close-world Datasets

- DARPA TC dataset
- In-lab Arena dataset
- NewlySim dataset<sup>1</sup>: contains attack behaviors that happened after CTI reports are released.

## • Open-world Dataset

- An open-world dataset across over 180 endpoints

## • Metrics

- graph-level precision/recall, node-level precision/recall

## • Baselines

- Detection baselines: HOLMES, AIRTAG, NodLink, KAIROS, EXTRACTOR+POIROT
- CTI extraction baselines: TPPDrill, LADDER, EXTRACTOR

TABLE IV: Summary of our evaluation datasets. “Duration” and “Event Rate” denote the duration of data collection and the average number of events generated per second, respectively.

Dataset	# APTs	Duration	# Hosts	Event Rate	# Attack Actions
THEIA	1	247h	1	11.25 eps	97
TRACE	2	264h	1	75.76 eps	93
In-lab Arena	5	144h	5	48.23 eps	202
NewlySim	2	336h	3	168.69 eps	39
Open-World	6	168h	186	28.13eps	212

<sup>1</sup><https://github.com/myh0301/KNOWHOW>

# Evaluation: Effectiveness

- **Graph-level accuracy:**
  - Detects **all the attacks** and only reports **0 false positives**.
- **Node-level accuracy:**
  - Node-level precision: KnowHow reduces up to **90% more node-level false positives** compared to existing detection baselines.
  - Node-level recall: KnowHow achieves almost perfect recall.

TABLE VI: The graph-level accuracy results for KNOWHOW and baselines. P stands for precision, and R stands for recall.

Detection Baselines	Dataset									
	THEIA		TRACE		In-Lab Arena		NewlySim		Open-World	
	P	R	P	R	P	R	P	R	P	R
<b>HOLMES</b>	<b>1.00</b>	<b>1.00</b>	0.15	<b>1.00</b>	0.04	<b>1.00</b>	0.14	<b>1.00</b>	0.15	0.40
<b>AIRTAG</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	0.83	<b>1.00</b>	1.00	<b>1.00</b>	0.63	<b>1.00</b>
<b>NODLINK</b>	<b>1.00</b>	<b>1.00</b>	0.67	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	0.71	<b>1.00</b>
<b>KAIROS</b>	0.91	<b>1.00</b>	0.88	0.88	<b>1.00</b>	<b>1.00</b>	0.50	<b>1.00</b>	0.63	<b>1.00</b>
<b>EXTRACTOR + POIROT</b>	0.33	<b>1.00</b>	0.33	<b>1.00</b>	0.50	<b>1.00</b>	0.25	<b>1.00</b>	0.43	0.30
<b>KNOWHOW</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>	<b>0.91</b>	<b>1.00</b>

TABLE VII: The node-level accuracy results for KNOWHOW and baselines. P stands for precision, and R stands for recall.

Detection Baselines	Dataset									
	THEIA		TRACE		In-Lab Arena		NewlySim		Open-World	
	P	R	P	R	P	R	P	R	P	R
<b>HOLMES</b>	0.01	0.98	0.01	0.74	0.01	0.32	0.01	0.40	0.04	0.21
<b>AIRTAG</b>	0.31	0.84	0.26	0.88	0.18	0.96	0.19	0.86	0.30	0.87
<b>NODLINK</b>	0.23	<b>1.00</b>	0.25	<b>0.98</b>	0.17	0.92	0.28	0.94	0.48	0.90
<b>KAIROS</b>	0.13	0.93	0.11	0.94	0.32	0.92	0.17	0.96	0.33	0.94
<b>EXTRACTOR + POIROT</b>	0.34	0.77	0.34	0.88	0.56	0.54	0.29	0.25	0.38	0.22
<b>KNOWHOW</b>	<b>0.62</b>	<b>1.00</b>	<b>0.82</b>	<b>0.98</b>	<b>0.82</b>	<b>1.00</b>	<b>0.78</b>	<b>1.00</b>	<b>0.82</b>	<b>1.00</b>

# Evaluation: Effectiveness

- **Effectiveness of gIoC:**

- KnowHow can identify attack behaviors in different datasets using gIoC, while failing when using results from the other SOTA CTI extraction.
  - Because, using the event semantic lifting method, KnowHow can detect the malicious process and file node **from the granularity of behavior** and maps the gIoC, dealing with the situation when the malicious file is **renamed to disguise** itself as a normal file.

TABLE VIII: The graph-level accuracy results for KNOWHOW and baselines. P stands for precision, and R stands for recall.

Dataset	CTI Extraction Baselines						Ours	
	TTPDRILL + KNOWHOW *		LADDER + KNOWHOW *		KNOWHOW with only IOC		KNOWHOW	
	P	R	P	R	P	R	P	R
THEIA	0.17	<b>1.00</b>	0.20	<b>1.00</b>	0.50	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
TRACE	0.50	<b>1.00</b>	0.40	<b>1.00</b>	0.67	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
In-Lab Arena	0.56	<b>1.00</b>	0.33	0.80	0.83	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
NewlySim	0.17	<b>1.00</b>	0.17	<b>1.00</b>	0.17	<b>1.00</b>	<b>1.00</b>	<b>1.00</b>
Open-World	0.67	0.80	0.38	0.30	0.35	0.60	<b>0.91</b>	<b>1.00</b>

TABLE IX: The node-level accuracy results for KNOWHOW and baselines. P stands for precision, and R stands for recall.

Dataset	CTI Extraction Baselines						Ours	
	TTPDRILL + KNOWHOW *		LADDER + KNOWHOW *		KNOWHOW with only IOC		KNOWHOW	
	P	R	P	R	P	R	P	R
THEIA	0.31	<b>1.00</b>	0.15	0.54	0.23	0.88	<b>0.62</b>	<b>1.00</b>
TRACE	0.67	<b>0.98</b>	0.12	0.63	0.18	0.83	<b>0.82</b>	<b>0.98</b>
In-Lab Arena	0.54	0.92	0.21	0.52	0.32	0.65	<b>0.82</b>	<b>1.00</b>
NewlySim	0.12	0.84	0.08	0.12	0.19	0.68	<b>0.78</b>	<b>1.00</b>
Open-World	0.59	0.37	0.11	0.37	0.11	0.27	<b>0.82</b>	<b>1.00</b>

# Evaluation: Technique Label Accuracy

- Our datasets cover 643 attack actions spanning 12 tactics and 65 techniques.
- KnowHow accurately labels 559 of all 643 attack actions, accounting for **87.0% accuracy**.
- After manual inspection, **507 out of 559 actions can be labeled by gIoCs** while only 101 actions can also be identified by IoCs, showing the effectiveness of gIoCs.

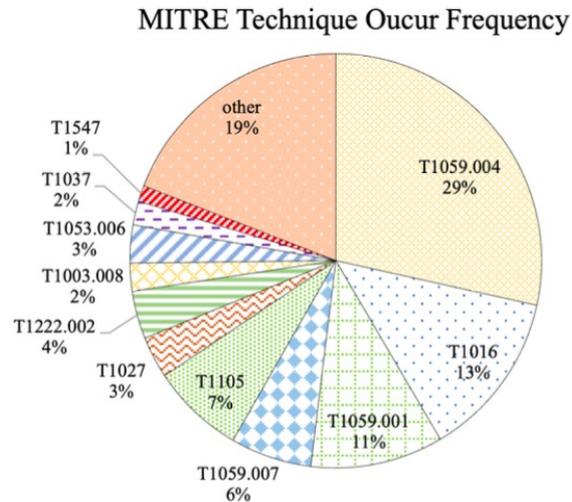


Fig. 5: Top 10 frequent techniques in our experiments.

# Evaluation: Efficiency

- **Throughput:** how many system events can be processed per second (eps)
  - KnowHow is **comparable** to SOTA online detection systems.
  - KnowHow is capable of monitoring over **180** hosts online in our real-world Open-World Dataset.

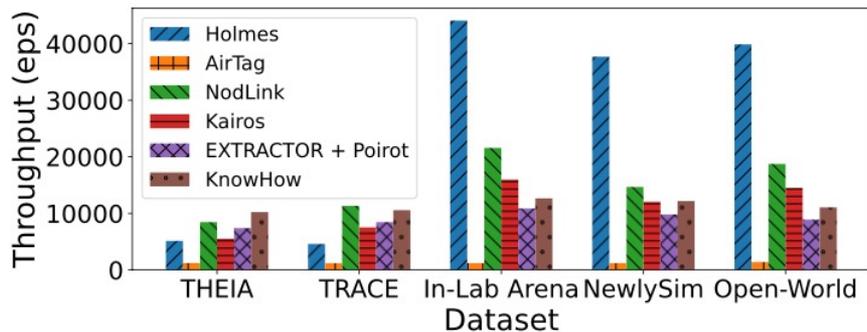


Fig. 6: Working throughput among different frameworks. TTPDRILL and LADDER have the same throughput as KNOWHOW.

- **Mimicry Attacks**

- We constructed a mimicry dataset, Mimic-Prov, following the steps described in the previous work, and evaluated KnowHow on it.
- KnowHow can successfully detect **all attack** events at the graph level across **all these mimicry insertion ratios** within the Mimic-Prov dataset.

- **Incomplete Attacks**

- We simulate incomplete attacks by removing attack steps from the original attack sequences in the In-lab Arena dataset.
- Both graph- and node-level detection performance of KnowHow **remain stable and effective**, except for a limited increase in node-level false positives, when the attack integrity ratio is 0.80 or lower.

# Evaluation: Case Study on Unseen Attacks

- KnowHow successfully detects two unseen multi-stage APT attacks in NewlySim dataset.
- These attacks exploited **newly disclosed** CVEs that were **not yet included** in any CTI report used to build KnowHow's CKD, and **occurred after** those CTIs, ensuring they were truly unseen.

## KNOWHOW APT Attack Report

### Incident Overview

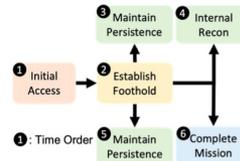
Date: 2024-05-10  
Reported By: KNOWHOW  
Incident ID: th-008

### Executive Summary

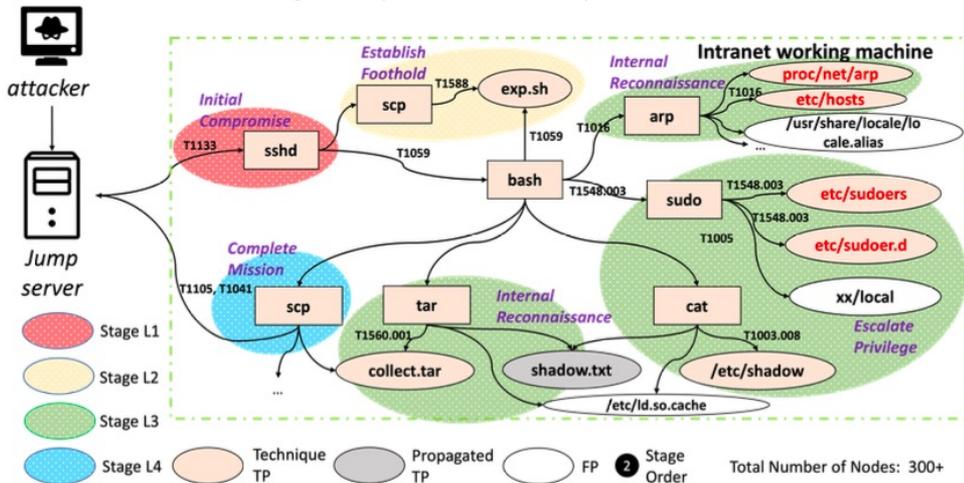
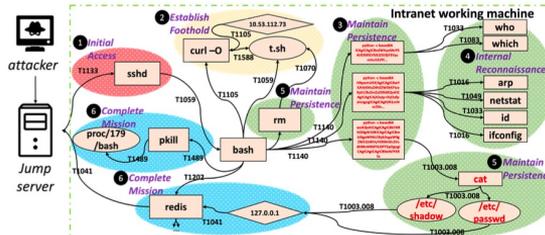
On 2024-05-10, KNOWHOW identified a sophisticated cyberattack originating from an internal IP address (192.0.1.15). The attack involved multiple stages including initial access, establishing a foothold, maintaining persistence, internal reconnaissance, and completing the mission. The attacker utilized various techniques and tools to achieve their objectives.

### Attack Overview

#### 1. Attack Lifecycle Path



#### 2. Attack Provenance Graph



# Conclusion

---

- We propose KnowHow, a CTI-knowledge-driven online provenance analysis solution that can **automatically apply high-level attack knowledge** from CTI reports to detect APT attacks in low-level system events.
- Our experiments show that KnowHow outperforms existing baselines in terms of both **accuracy** and **interpretability** while maintaining a **comparable throughput**. Moreover, KnowHow can achieve great **robustness** to attacks stemming from unknown vulnerabilities and existing mimicry attacks.



<https://github.com/myho301/KNOWHOW>

mengyuhan@pku.edu.cn