

# Device Tracking via Linux's New TCP Source Port Selection Algorithm\*

Moshe Kol, Amit Klein, Yossi Gilad

Hebrew University of Jerusalem

DANSS 2022-2023

\* Accepted to USENIX Security 2023

# Web-based User Tracking

- Web-based user (device) tracking is used for various purposes:
  - ✗ Real-time targeted marketing
  - ✗ Surveillance purposes
  - ✓ Fraud detection
  - ✓ Campaign measurement
  - ✓ Protection against account hijacking



Data mining  
Anti-bot services  
Enterprise security management  
Limiting number of accesses to services

# Third-Party Cookies

- Set by **third-party content** (Ads, Social media, etc.)
- Traditionally (ab)used for **cross-site tracking**
- **Slowly dying**
  - Unsuitable for **cross-browser** tracking
  - Users can **delete** them
  - Under **regulation** (GDPR, CCPA)
  - **Browser support is phased-out**



⇒ **Trackers are now looking for new tracking technologies**

# Our Work

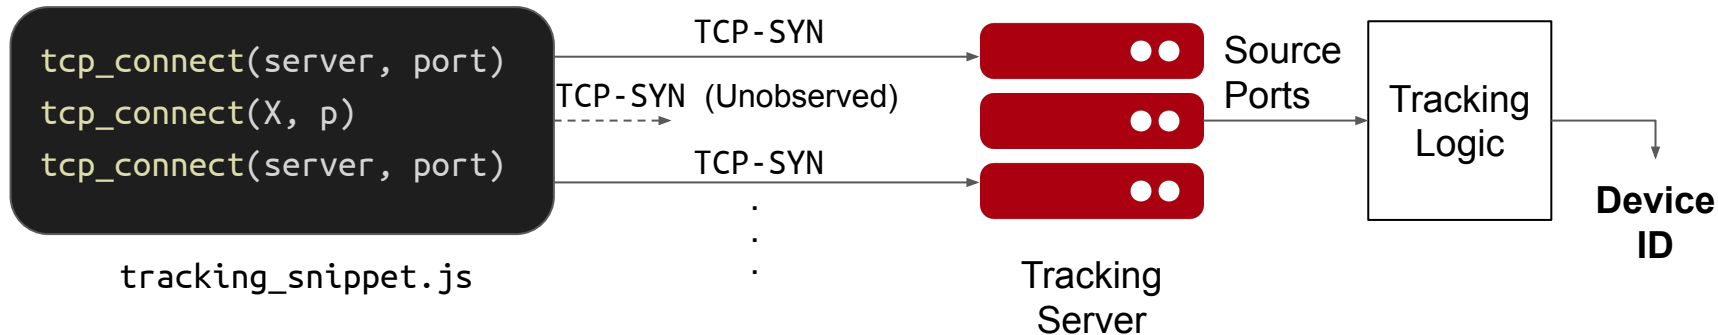
- Our work: Uses **TCP source ports** for user tracking
- Entry point: A **tracking snippet** (JavaScript)
- The snippet runs on **the user browser**, executes the **tracking logic** and obtains a **device ID**



# Our Work


TCP source ports are not exposed via JavaScript  $\Rightarrow$  **tracking server** is used:

- The **tracking snippet** forces the user device to establish TCP connections
- The **tracking server** analyzes source port measurements for the device ID



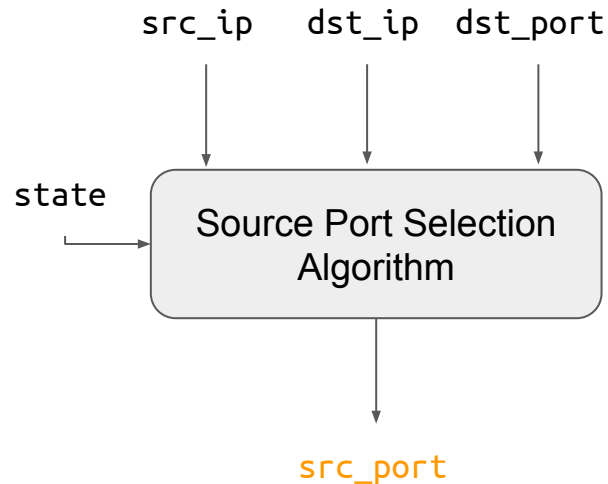
# TCP Source Port Selection

- TCP connections are identified by the **4-tuple**:  
(src\_ip, src\_port, dst\_ip, dst\_port)
- Client provides (dst\_ip, dst\_port) at **connect()**
  - src\_ip is automatically set by the OS
- OS chooses src\_port (“Ephemeral Port”)



# TCP Source Port Selection

- Source ports should be **unpredictable** and **infrequently reused**
- **RFC 6056** defines 5 port selection algorithms
- Linux moved to the **Double-Hash Port Selection (DHPS) Algorithm** starting from kernel 5.12



# RFC 6056, Algorithm 4: Double-Hash Port Selection (DHPS)

- Algorithm state:
  - `table`: Table of 32-bit integers of length  $T$  (initialized randomly at boot time)
  - $K_1, K_2$ : Two 128-bit keys (initialized randomly at boot time)
  - $F, G$ : Two cryptographic keyed hash functions
- To select `src_port` for a given 3-tuple (`src_ip`, `dst_ip`, `dst_port`):
  - Compute `offset` =  $F(K_1, \text{src\_ip}, \text{dst\_ip}, \text{dst\_port}) \in [0, 2^{32}-1]$
  - Compute `index` =  $G(K_2, \text{src\_ip}, \text{dst\_ip}, \text{dst\_port}) \in [0, T-1]$
  - Compute `port` =  $\text{min\_ephemeral} + (\text{table}[\text{index}] + \text{offset}) \% \text{num\_ephemeral}$
  - Increment `table[index]` by 1



# RFC 6056, Algorithm 4: Security Issues

- Issue #1: Shared global state
  - Across network interfaces, protocols and network namespaces
- Issue #2: Small table length
  - RFC 6056 suggests  $T = 10$ ; Linux uses  $T = 256$
- Issue #3: Deterministic change of state
  - Table cells are incremented by 1



## Device Tracking based on DHPS

- Algorithm state is **shared** between **Internet-facing** and **loopback** interfaces
- Loopback has a **fixed** IP address: **127.0.0.1**

TCP-Connect(“127.0.0.1”, **x**)

⇒ **index** = G(**K<sub>2</sub>**, “127.0.0.1”, “127.0.0.1”, **x**)

# Device Tracking based on DHPS

- The idea: Collect **hash collisions** of loopback traffic

i.e. Pairs ( $x$ ,  $y$ ) such that

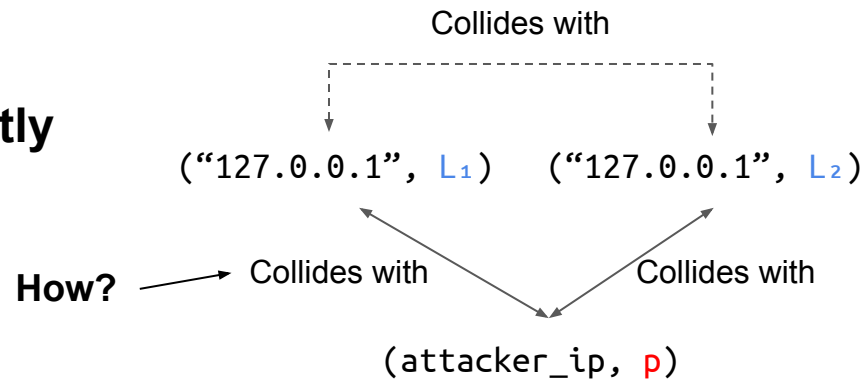
$$G(K_2, "127.0.0.1", "127.0.0.1", x) = G(K_2, "127.0.0.1", "127.0.0.1", y)$$

- The pairs  $\{(x_i, y_i)\}$  are **network independent** (depend only on  $K_2$ )  
 $\Rightarrow \{(x_i, y_i)\}$  form a **device ID**.
- Device ID persists across browsers, network switches, etc.
  - Lasts as long as  $K_2$  persists (until reboot, in Linux)

# Device Tracking based on DHPS

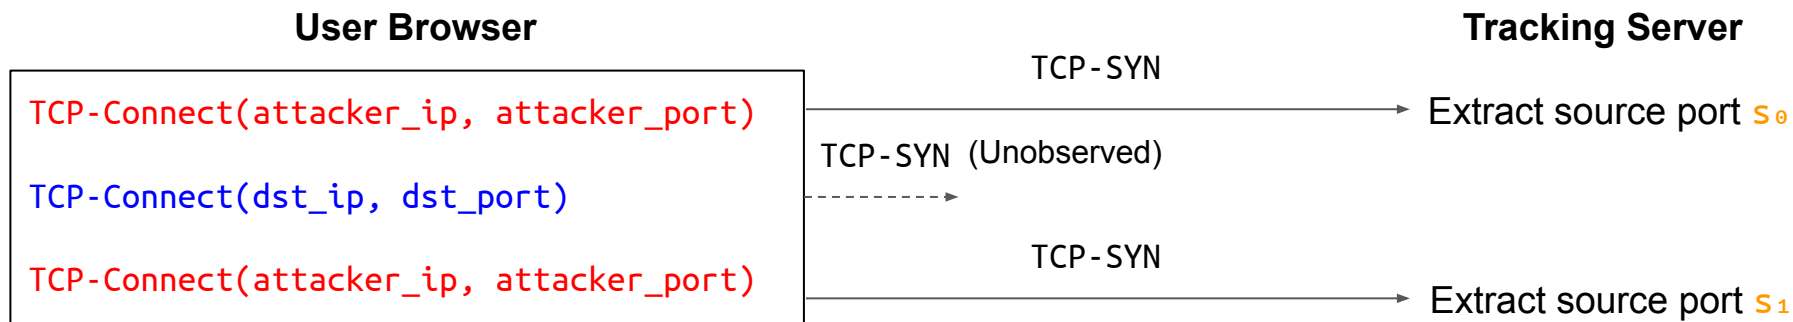
- Challenge: No access to source port information from JavaScript
  - Also, no access to the algorithm state (`table`,  $K_1$ ,  $K_2$ )
- We can only observe the source ports directed to the attacker server

- We will find loopback collisions **indirectly**



# Device Tracking based on DHPs: The Primitive

The user's device is forced (via JS) to establish 3 TCP connections:



Attacker computes the source port difference  $\Delta = S_1 - S_0$ .


If there was a collision then  $\Delta = 2$  (the table cell was incremented twice).

Otherwise,  $\Delta = 1$ .

# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 1: Collect  $T$  attacker destinations that uniquely cover all `table` cells




# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 1: Collect  $T$  attacker destinations that uniquely cover all `table` cells

The attacker **doesn't know** the mapping, only that it **exists**.

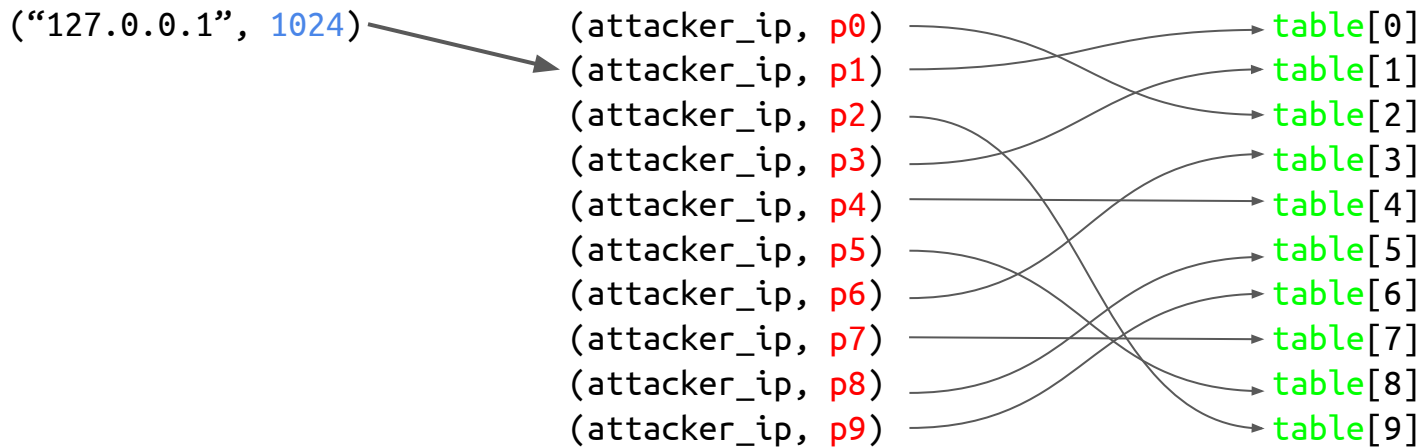


# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions

Fingerprint:



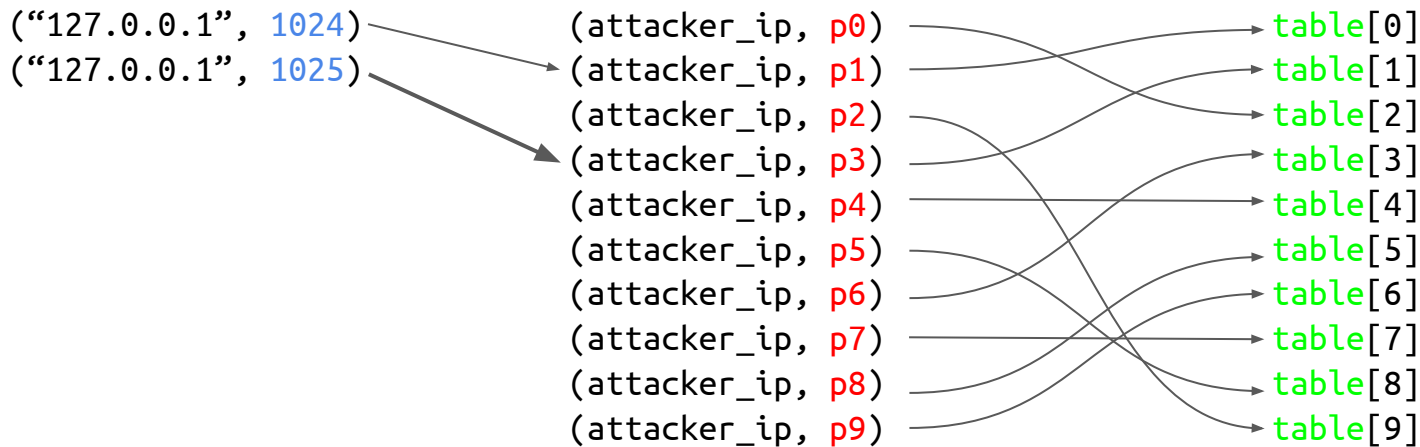


# Device Tracking based on DHCP: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions

Fingerprint:

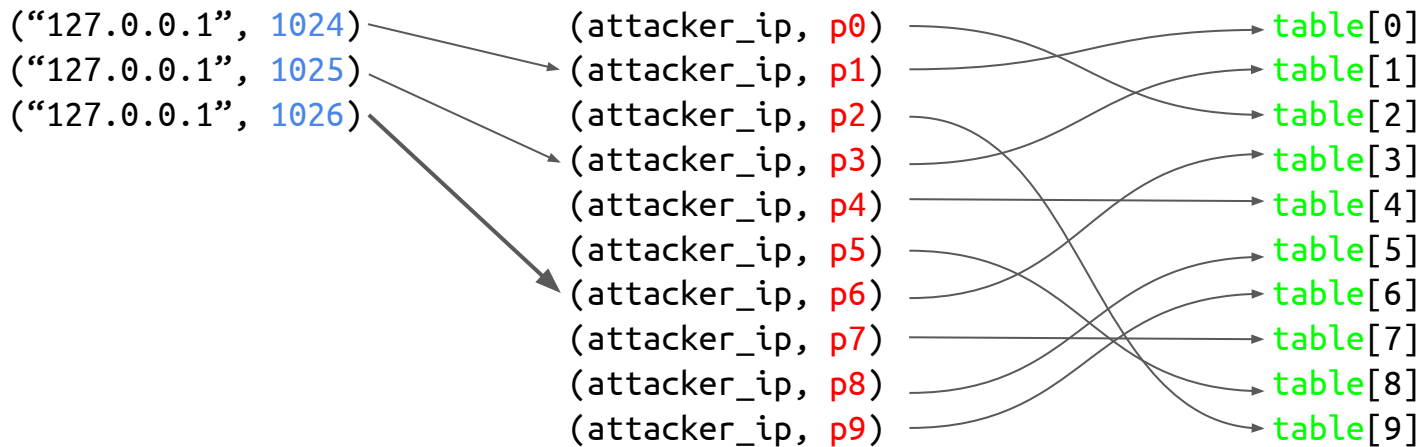


# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions

Fingerprint:

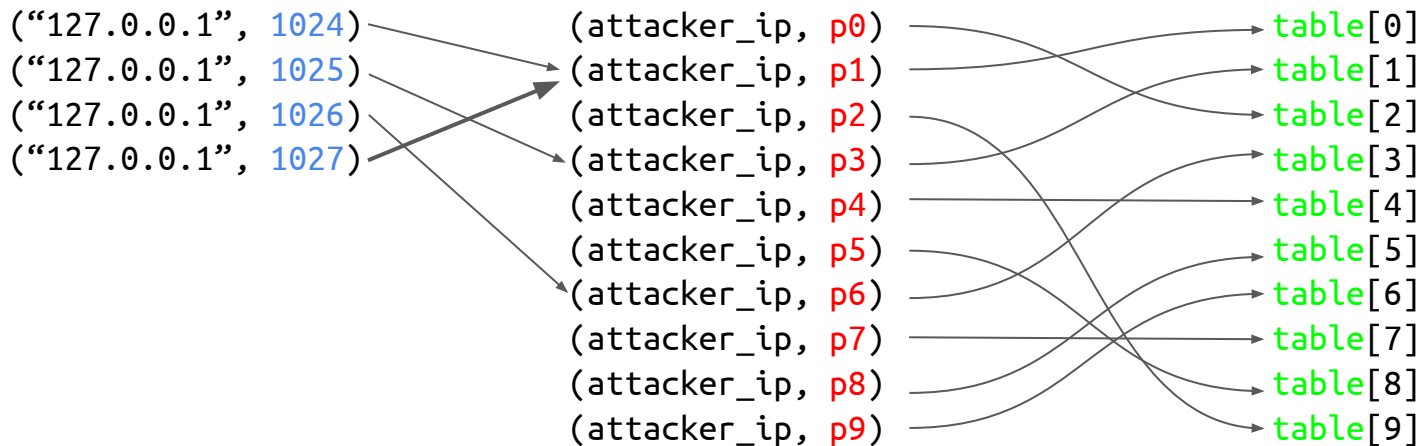


# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions

Fingerprint:  
(1024, 1027)

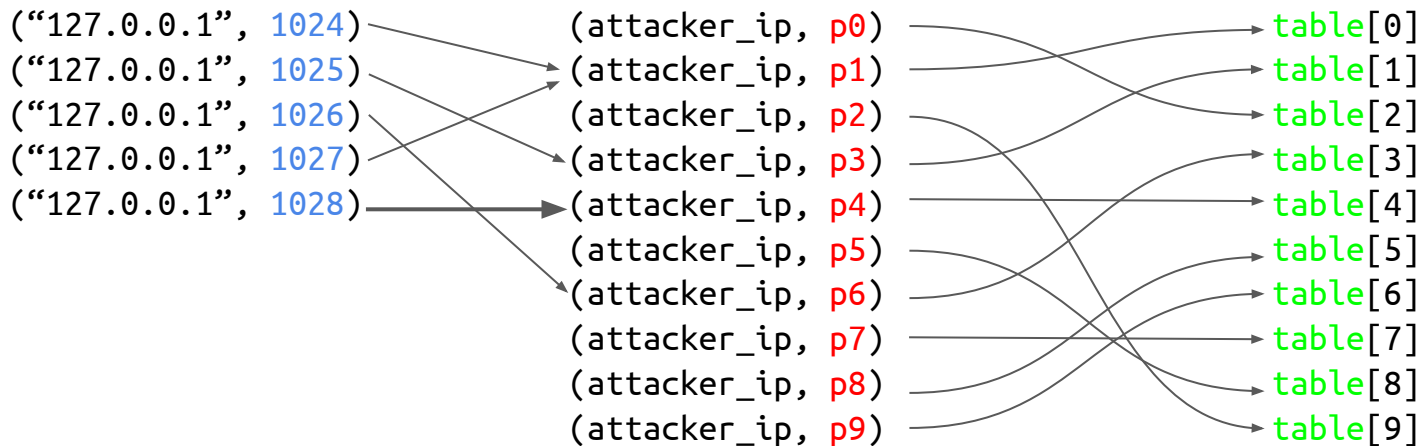


# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions

Fingerprint:  
(1024, 1027)

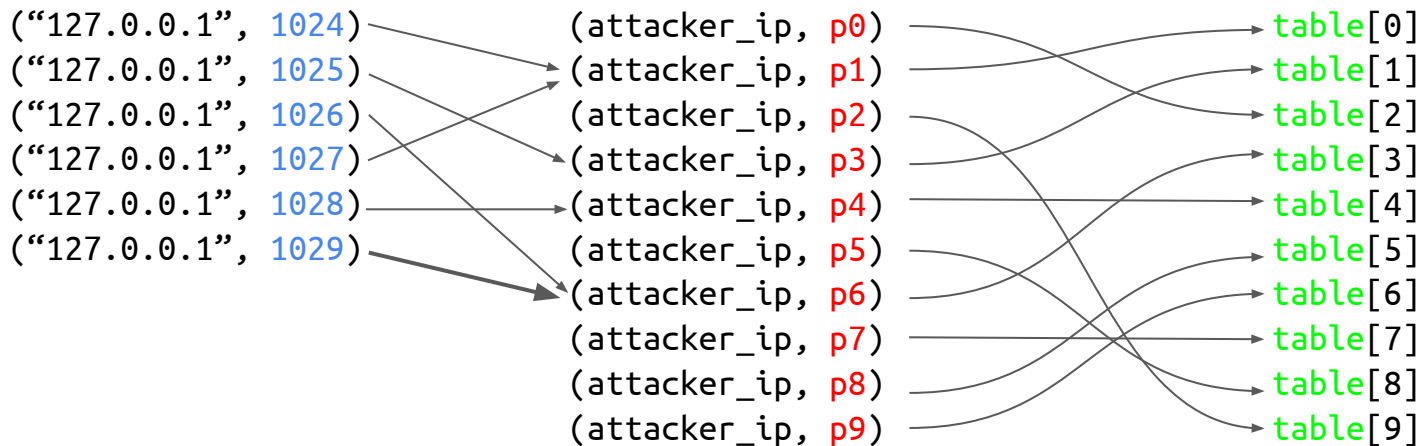


# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions

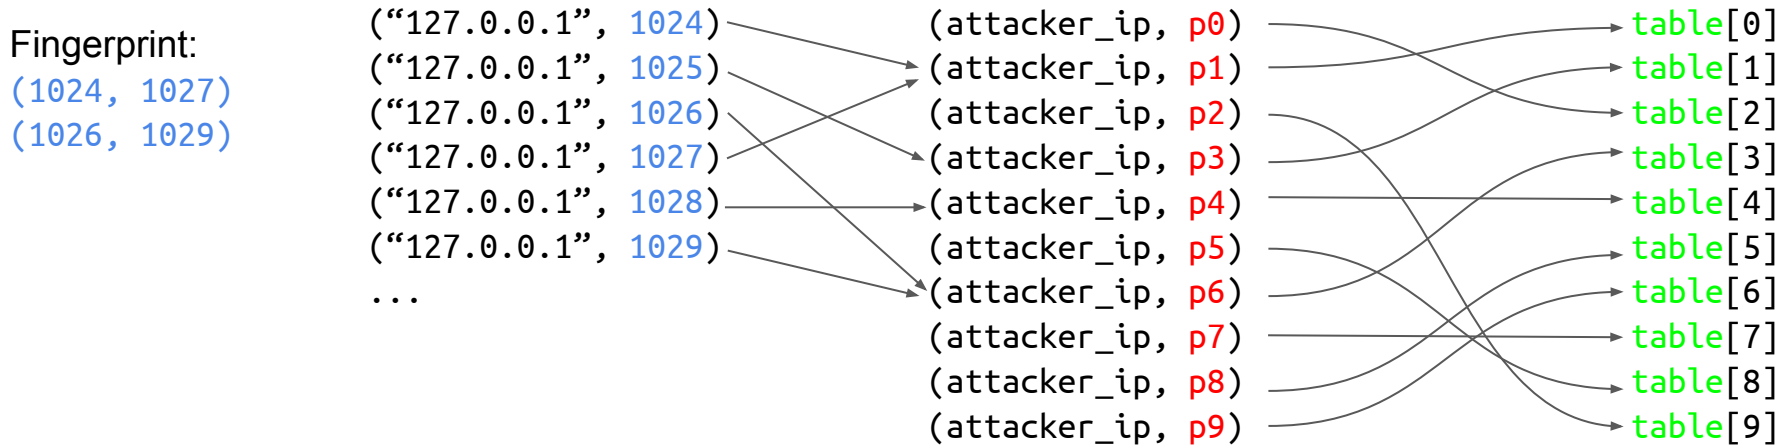
Fingerprint:  
(1024, 1027)  
(1026, 1029)



# Device Tracking based on DHPS: The Plan

Our attack works in 2 phases:

- Phase 2: Collect loopback hash collisions



## Device Tracking based on DHPS: Phase 1

- The Goal: Obtain  $T$  attacker destinations that uniquely cover every **table** cell
- The algorithm is **iterative**:
  - $U$  := set of **unique** attacker destinations
  - We begin with  $U = \emptyset$
  - Expand  $U$  at every iteration, until  $|U| = T$  (every table cell is covered)



# Device Tracking based on DHPS: Phase 1

- On each iteration: Extend  $U$  with more unique attacker destinations.
  - Let  $S$  be a fresh **set of  $T$**  random attacker destinations.\*
  - Force the user browser to execute (in-order):

TCP-Connect( $S$ )

TCP-Connect( $U$ )

TCP-Connect( $S$ )

- Update  $U$  with all  $a \in S$  such that  $\Delta_a = 1$ .

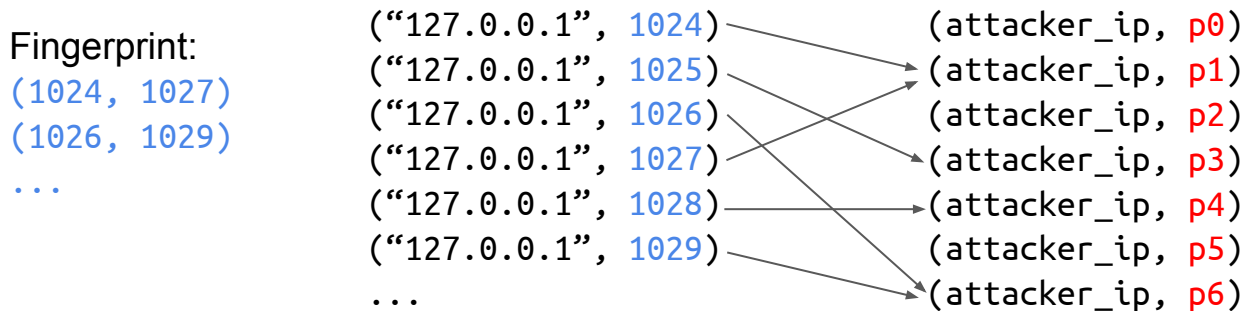
S:	$a_0$	$a_1$	$a_2$	$a_3$	$a_4$	$a_5$	$a_6$	$a_7$	$a_8$
$\Delta$ :	2	3	1	1	2	2	1	3	2

\*  $T$  maximizes the expected number of new unique attacker destinations.



## Device Tracking based on DHCP: Phase 2

- The Goal: Collect loopback collisions
- We do this indirectly:
  - We first map loopback destinations to attacker destinations
  - Then, we infer loopback collisions from the mapping



## Device Tracking based on DHPS: Phase 2

- On each iteration  $\ell$ :

Find the attacker destination that collides with (“127.0.0.1”,  $\ell+1024$ ):

TCP-Connect(U)

TCP-Connect(“127.0.0.1”,  $\ell+1024$ )

TCP-Connect(U)

⇒ Look for  $a \in U$  s.t.  $\Delta_a > 1$ .

U:	$u_0$	$u_1$	$u_2$	$u_3$	$u_4$	$u_5$	$u_6$	$u_7$	$u_8$
$\Delta$ :	1	1	1	2	1	1	1	1	1

- We stop when “enough” loopback collisions were collected

We guarantee that:  $\Pr[\mathbf{sig}(D) = \mathbf{sig}(D')] \leq \frac{1}{\binom{N}{2}}$

$N$ : population size  
 $D, D'$ : random devices

## Device Tracking based on DHPs: Improvements

- Grouping loopbacks: Test  $\alpha$  loopbacks instead of one, in phase 2.

TCP-Connect(U)

TCP-Connect("127.0.0.1", 1024)  $\times 2^0$

TCP-Connect("127.0.0.1", 1025)  $\times 2^1$

TCP-Connect("127.0.0.1", 1026)  $\times 2^2$

TCP-Connect("127.0.0.1", 1027)  $\times 2^3$

TCP-Connect(U)

Phase 2 ( $\alpha=4$ )

# Device Tracking based on DHPS: Improvements

- Robustify against organic noise: Tolerate up to  $\beta$  noise.

TCP-Connect(U)

TCP-Connect("127.0.0.1", 1024)  $\times 2^0 \times \beta$

TCP-Connect("127.0.0.1", 1025)  $\times 2^1 \times \beta$

TCP-Connect("127.0.0.1", 1026)  $\times 2^2 \times \beta$

TCP-Connect("127.0.0.1", 1027)  $\times 2^3 \times \beta$

TCP-Connect(U)

Phase 2 ( $\alpha=4$ )

**Note: Phase 1 already tolerates noise**

## Device Tracking based on DHCP: Limitations

- Our technique relies on **observing the machine-generated source ports**
- Cannot track proxy (Tor) users
- Cannot track under networks with port-rewriting NATs



# DHPS Implementation in Linux

- DHPS is implemented in `__inet_hash_connect()`
  - Used by both IPv4 and IPv6 code

- Table is **global** and its size is **256**

```
#define INET_TABLE_PERTURB_SHIFT 8
static u32 table_perturb[1 << INET_TABLE_PERTURB_SHIFT];
```

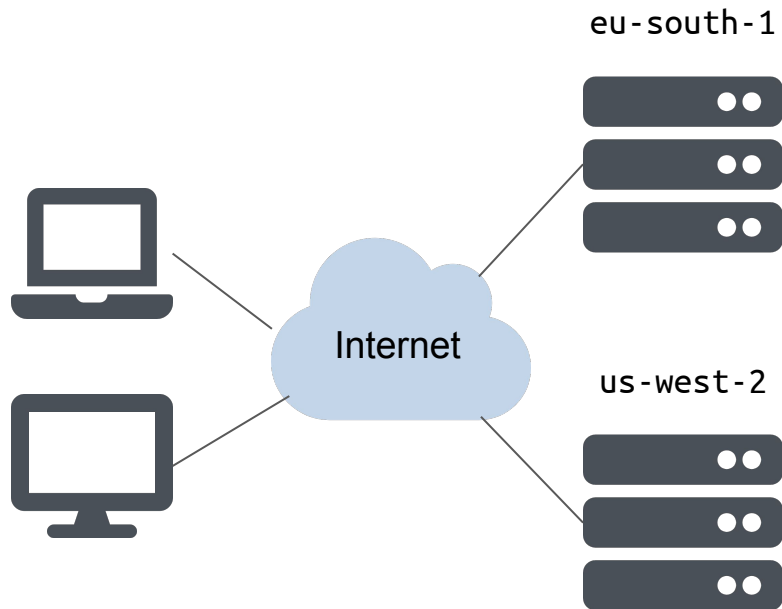
- SipHash with 128-bit key for F, G

```
u32 secure_ipv4_port_ephemeral(__be32 saddr, __be32 daddr, __be16 dport)
{
    net_secret_init();
    return siphash_3u32((__force u32)saddr, (__force u32)daddr,
                       (__force u16)dport, &net_secret);
}
```

- Noise Injection: Increment a table cell **twice** with probability 1/16

## Evaluation: Setup

- Prototype **tracking server** in Go and **tracking snippet** in HTML+JavaScript
- 2 tracking servers in different locations
- Multiple Linux laptops and PCs



# Evaluation: Results

## ✓ **Cross browser tracking**

- Tested Chrome v96.0 and Firefox v96





## Evaluation: Results

- ✓ Cross browser tracking
- ✓ **Cross browser privacy modes tracking**




## Evaluation: Results

- ✓ Cross browser tracking
- ✓ Cross browser privacy modes tracking
- ✓ **Dwell time: 5-15s on Chrome**



## Evaluation: Results

- ✓ Cross browser tracking
- ✓ Cross browser privacy modes tracking
- ✓ Dwell time: 5-15s on Chrome
- ✓ **Cross network tracking**
  - Consistency under IPv4 or IPv6
  - Works on some VPNs: TunnelBear and ExpressVPN



## Evaluation: Results

- ✓ Cross browser tracking
- ✓ Cross browser privacy modes tracking
- ✓ Dwell time: 5-15s on Chrome
- ✓ Cross network tracking
- ✓ **Android**
  - Tested Samsung Galaxy S21 with patched kernel
  - Dwell time: 18-21s on Chrome mobile



# Demo

The screenshot displays a terminal window with three main sections:

- Top Section:** A log of network traffic. It shows a sequence of 'loopback' events for ports 1888, 1889, and 1890. It then shows a 'connecting' event to 192.168.1.3 with 788 ports. This is followed by a 'phased' event with 'duration: 4.829990181s' and a 'phased\_iteration: 38'. A 'source port range' event is shown with '(configured): [127.0.0. 61044]' and '(observed): [20776, 80062]'. Finally, a 'Fingerprint' event is shown with '(1029, 1025): [(1030, 1035): (1036, 1039): (1409, 1417): (1488, 1478, 1476): (1485, 1484): (1485, 1486): (1485, 1486): (1485, 1486): (1485, 1486)]' and 'Fingerprint: fmd5: 02A03786'.
- Middle Section:** A terminal window titled 'RFC 6056 Demo Tracker'. It contains a form with the following fields:
  - Tracer address:
  - Refresh:
  - OS:
  - Refresh button:
- Bottom Section:** A log of network traffic, similar to the top section, showing 'loopback' events for ports 1888, 1889, and 1890. It then shows a 'connecting' event to 192.168.1.3 with 788 ports. This is followed by a 'phased' event with 'duration: 4.829990181s' and a 'phased\_iteration: 38'. A 'source port range' event is shown with '(configured): [127.0.0. 61044]' and '(observed): [20776, 80062]'. Finally, a 'Fingerprint' event is shown with '(1029, 1025): [(1030, 1035): (1036, 1039): (1409, 1417): (1488, 1478, 1476): (1485, 1484): (1485, 1486): (1485, 1486): (1485, 1486): (1485, 1486)]' and 'Fingerprint: fmd5: 02A03786'.

Full source code: <https://github.com/Oxkol/rfc6056-device-tracker>

## Vendor Status

- February 2022: Linux kernel security team was informed
- Assigned **CVE-2022-32296**
- Worked with the security team to patch the vulnerability
- May 2022: A patch was merged into the Linux kernel



# Countermeasures

<b>Countermeasure</b>	<b>Original Implementation</b>	<b>New Implementation</b>	<b>Effect</b>
<b>Increase table size</b>	T=256	T=64K	Collisions become less frequent

# Countermeasures

<b>Countermeasure</b>	<b>Original Implementation</b>	<b>New Implementation</b>	<b>Effect</b>
<b>Increase table size</b>	T=256	T=64K	Collisions become less frequent
<b>Periodic re-keying</b>	Every reboot	Every 10s	Collisions become useless after re-keying



# Countermeasures

Countermeasure	Original Implementation	New Implementation	Effect
<b>Increase table size</b>	T=256	T=64K	Collisions become less frequent
<b>Periodic re-keying</b>	Every reboot	Every 10s	Collisions become useless after re-keying
<b>More noise</b>	$\sim \text{Bernoulli}(1/16)$	$\sim \mathcal{U}\{0,7\}$	Collisions become harder to determine

## Conclusion

- DHPS (Algorithm 4 of RFC 6056) is vulnerable to device tracking
- Demonstrated our technique on Linux
- The device ID persist across browsers, browser privacy modes, networks, etc.
- Shows that user privacy can be undermined in non-obvious ways

# Thanks for Listening!

## Questions?

Extended paper: <https://arxiv.org/pdf/2209.12993.pdf>

Source code: <https://github.com/0xkol/rfc6056-device-tracker>

Demo video: <https://www.youtube.com/watch?v=pZbfV5nCQsA>