In the Room Where It Happens:
Characterizing Local Communication and Threats in Smart Homes

Aniketh Girish
IMDEA Networks Institute

Background

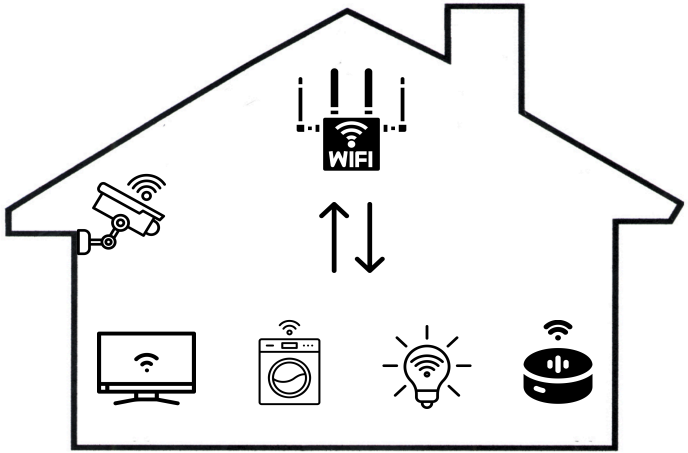

Background

Seamless integration and interoperability

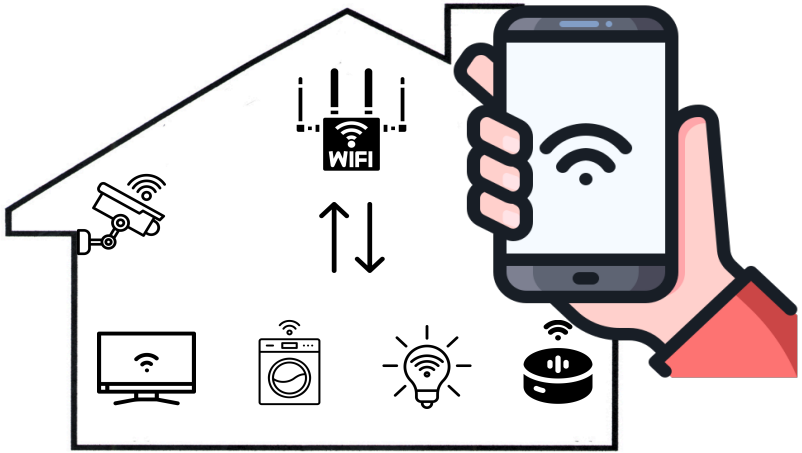

Background

Unicast traffic for command and control

Multicast/broadcast traffic for discovery

Seamless integration and interoperability

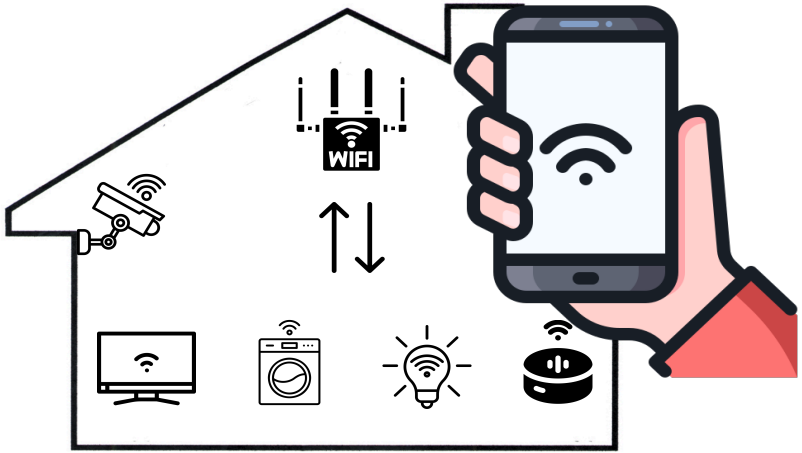

Background

Seamless integration and interoperability

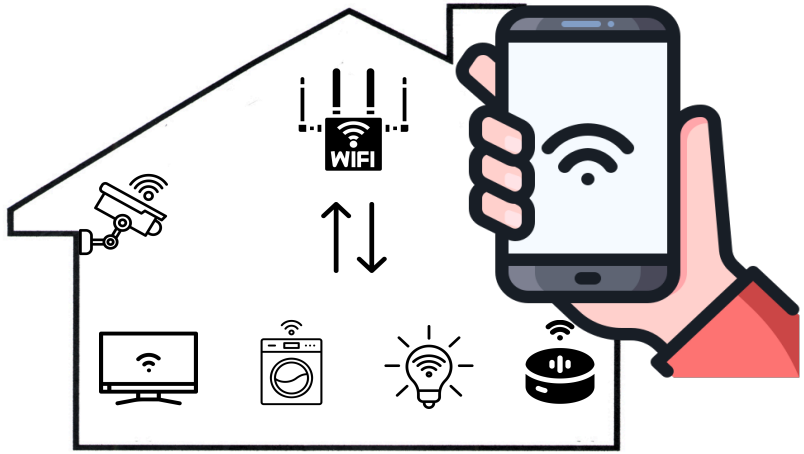

Unicast traffic for command and control

Multicast/broadcast traffic for discovery

Local communication and its associated threats are poorly understood

Prior work: study the devices or how IoT devices interact with cloud services

Background

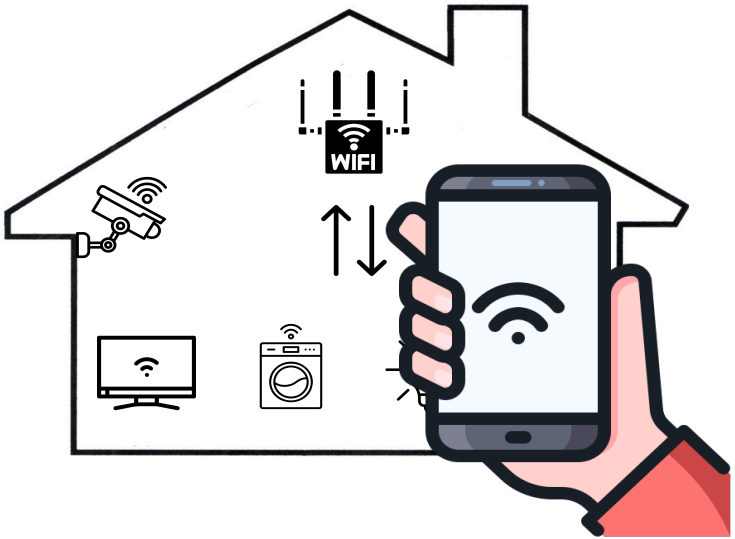

Background

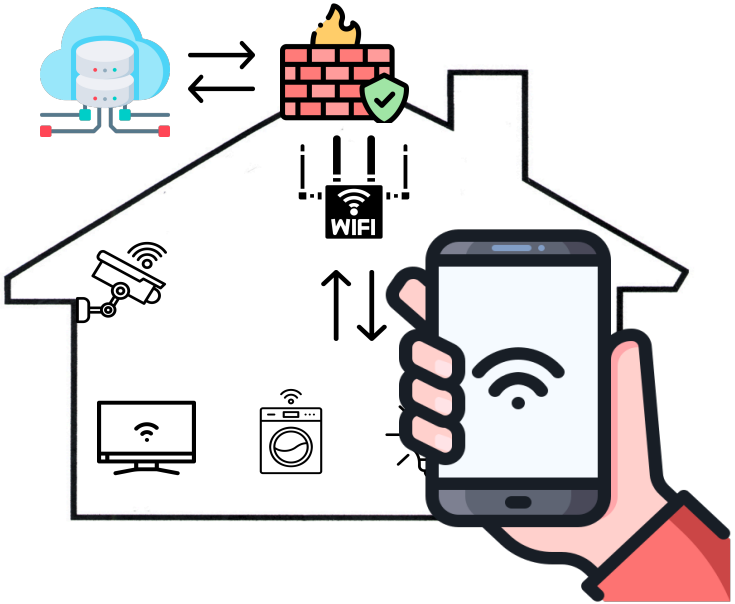

Background

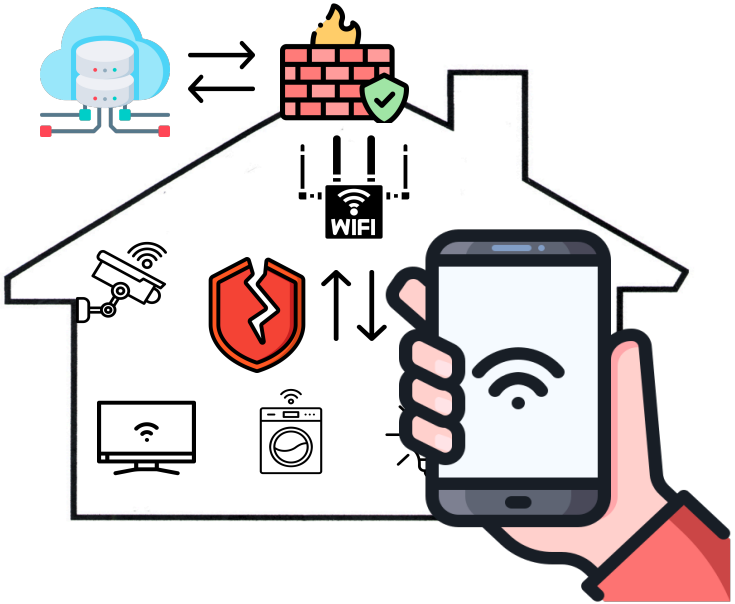

Broken local privacy protection

Background

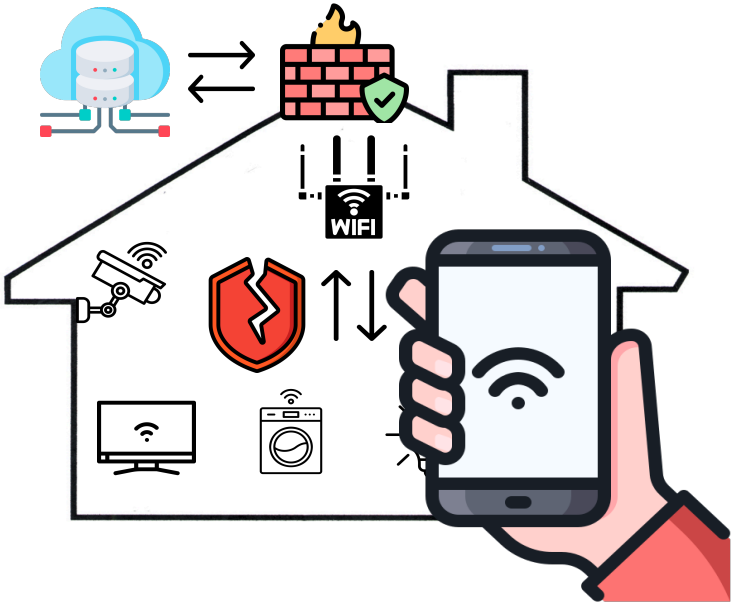

Broken local privacy protection

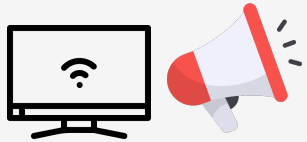

Device broadcast PII (MAC address, device IDs)

Background

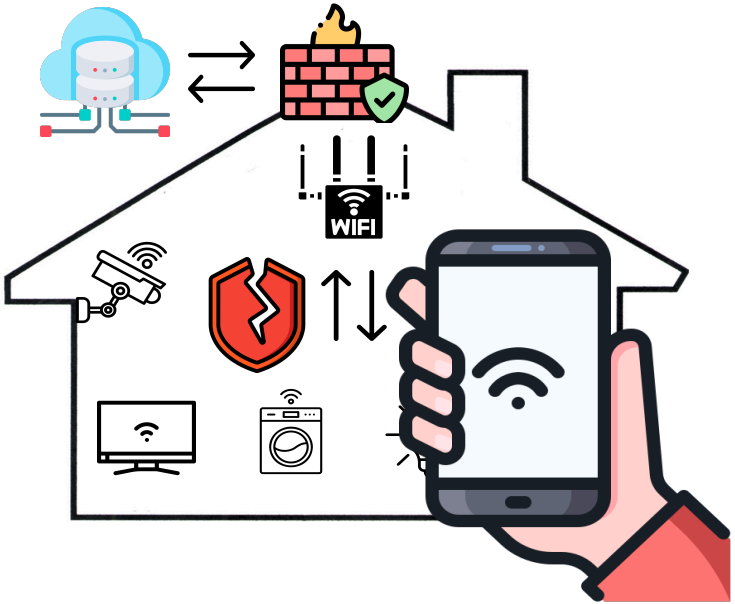

Broken local privacy protection

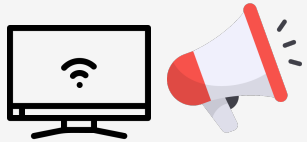

Device broadcast PII (MAC address, device IDs)

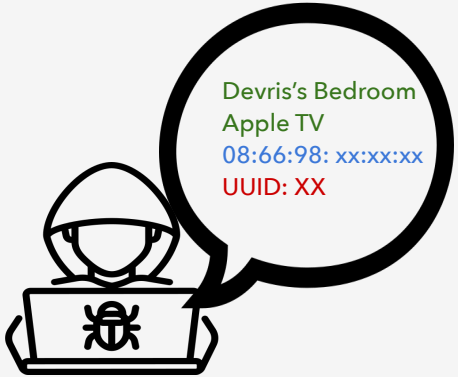

Background

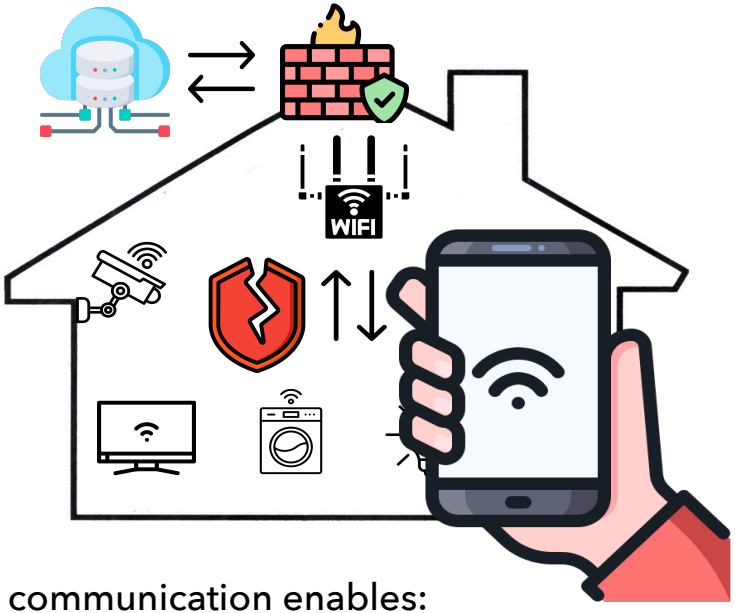

Local communication enables:

- **cross-device tracking**
- **unique household fingerprinting**
- **socio-economic status inference**

Broken local privacy protection

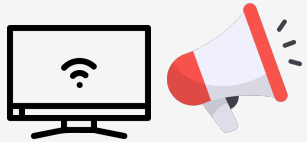

Device broadcast PII (MAC address, device IDs)

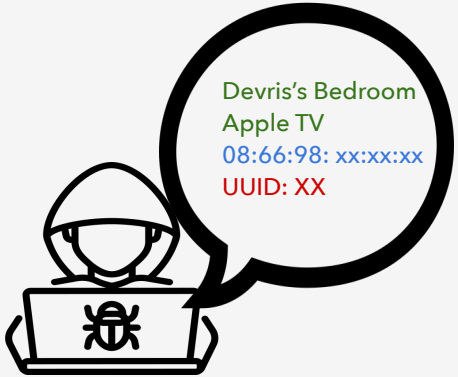

Surveillance & Tracking

Research Questions

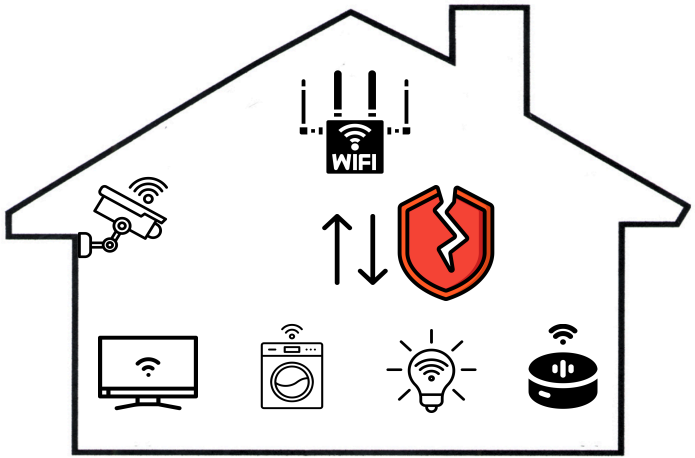

Research Questions

RQ1: What are the characteristics of smart home local network communication?

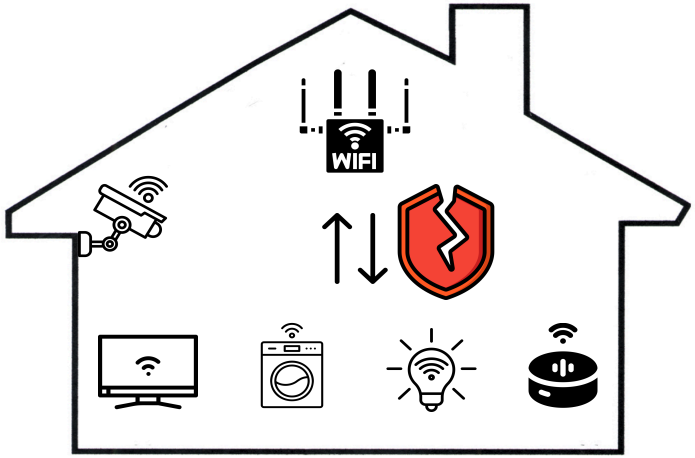

Research Questions

RQ1: What are the characteristics of smart home local network communication?

RQ2: What are the privacy and security threats?

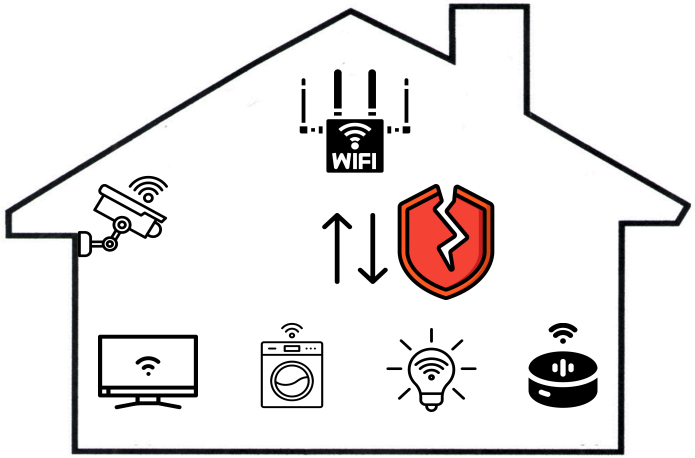

Research Questions

RQ1: What are the characteristics of smart home local network communication?

RQ2: What are the privacy and security threats?

RQ3: Is local network communication abused for fingerprinting and tracking?

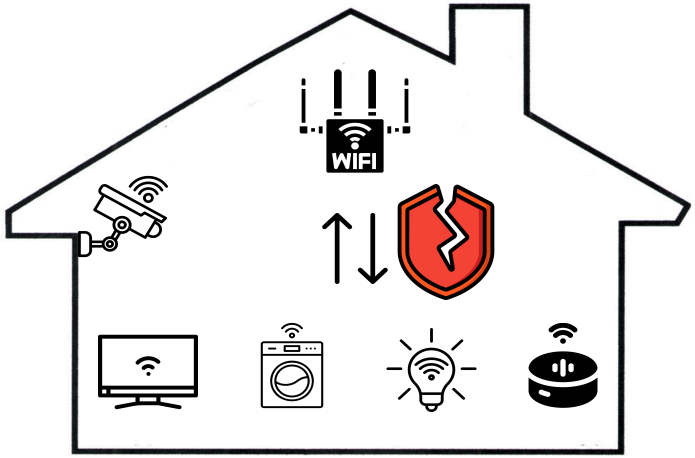

Our Testbed & Datasets

Devices: 93 consumer IP-based smart home devices.

Traffic: We capture all LAN traffic during interactions with IoT devices, and during idle periods.

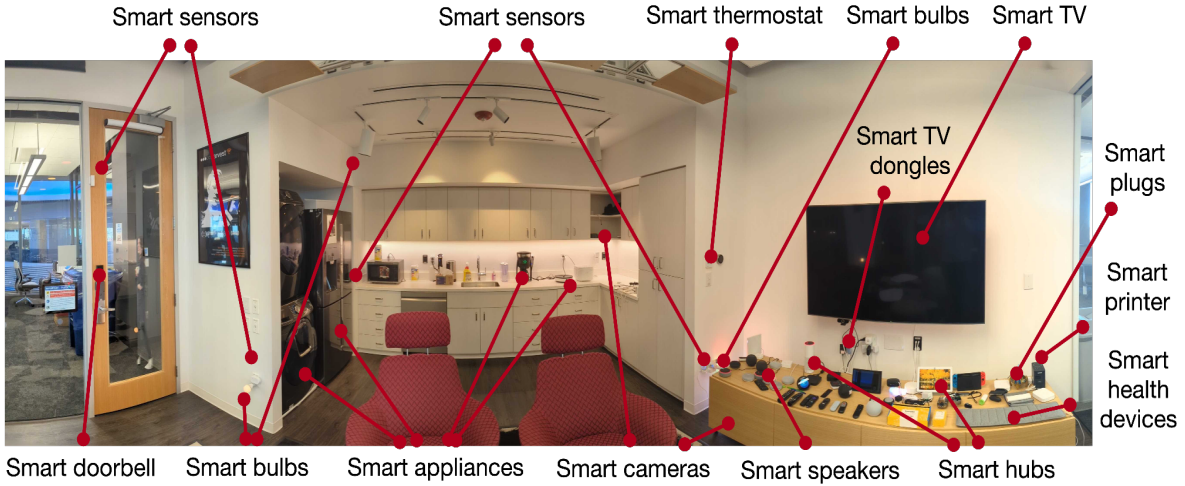

Our Testbed & Datasets

Devices: 93 consumer IP-based smart home devices.

Traffic: We capture all LAN traffic during interactions with IoT devices, and during idle periods.

Honeypot: Issues authentic responses to scan from IoT devices.

Active scan: nmap and Nessus.

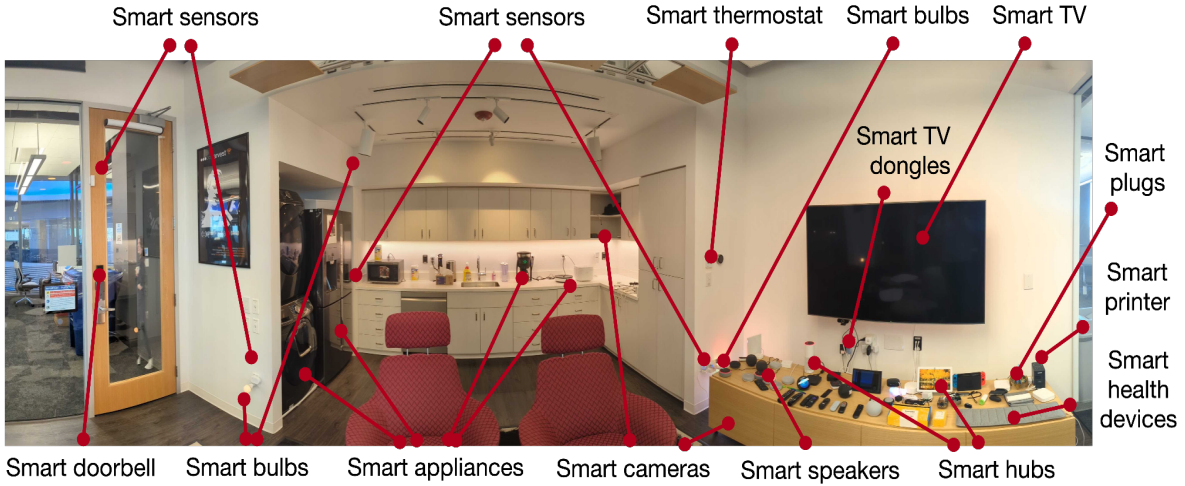

Our Testbed & Datasets

Devices: 93 consumer IP-based smart home devices.

Traffic: We capture all LAN traffic during interactions with IoT devices, and during idle periods.

Honeypot: Issues authentic responses to scan from IoT devices.

Active scan: nmap and Nessus.

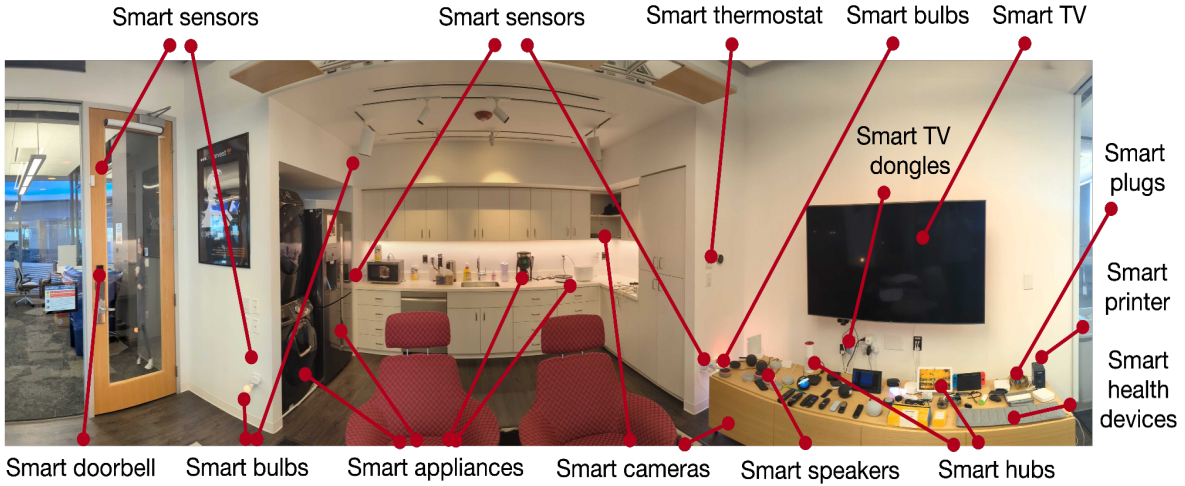

Our Testbed & Datasets

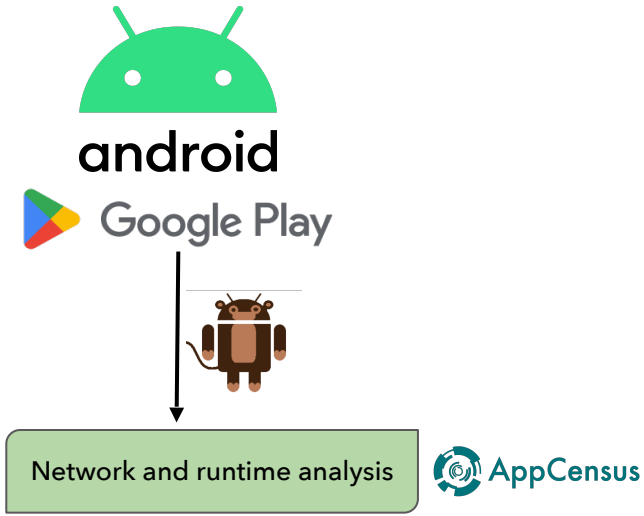

2,335 Android mobile apps:

- 987 IoT specific apps (e.g., companion apps).
- 1,348 randomly selected “regular” apps.

Our Testbed & Datasets

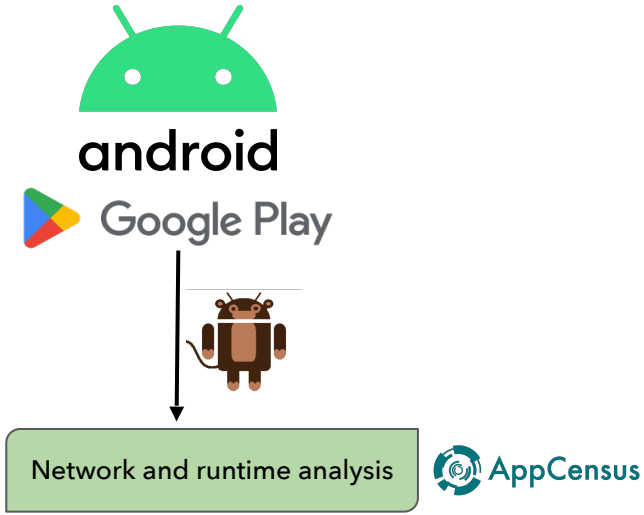

2,335 Android mobile apps:

- 987 IoT specific apps (e.g., companion apps).
- 1,348 randomly selected "regular" apps.

IoT Inspector

Crowdsourced IoT network traffic:

- 12,669 IoT devices from 3,860 households.
- 264 products from 165 vendors.
- mDNS and SSDP responses.

How do these devices interact with each other?

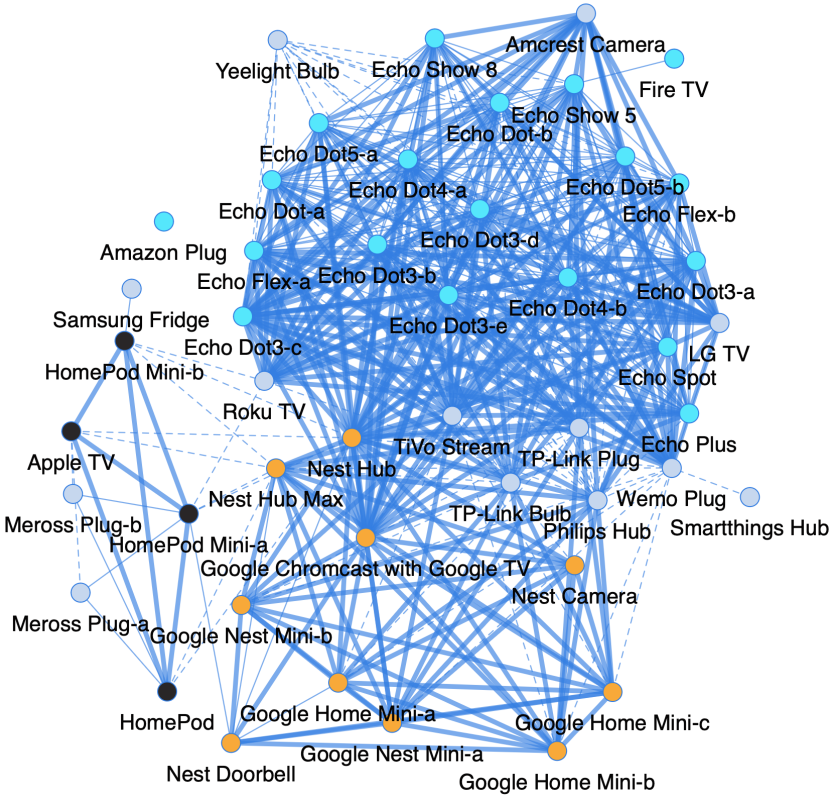

How do these devices interact with each other?

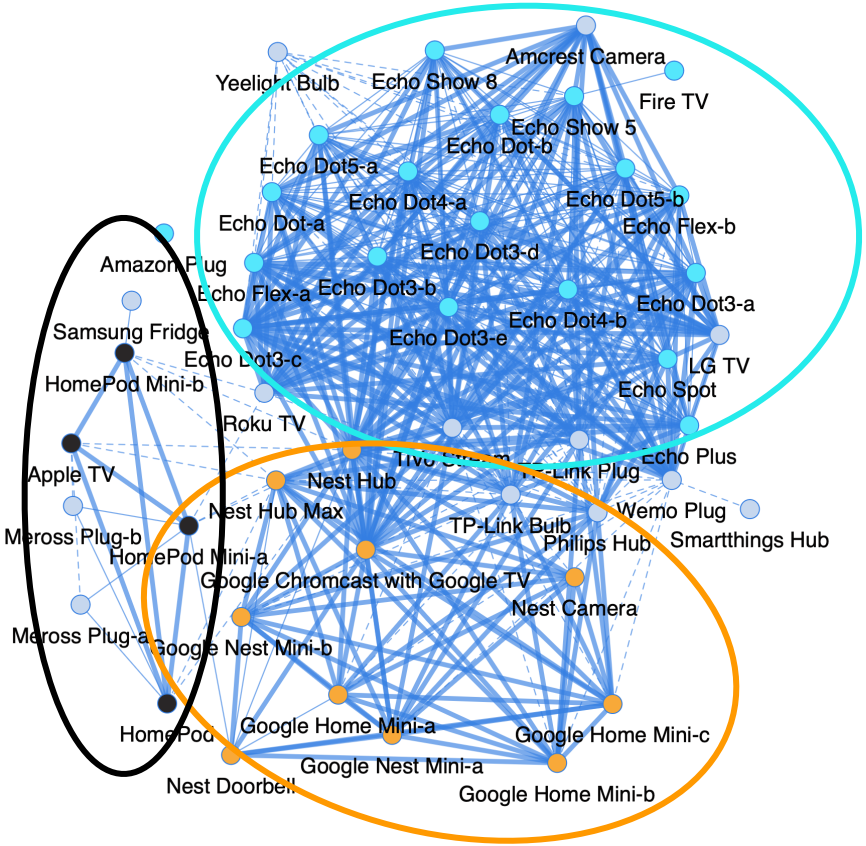

Intra-vendor communication across devices in Amazon, Google, and Apple's ecosystem.

How do these devices interact with each other?

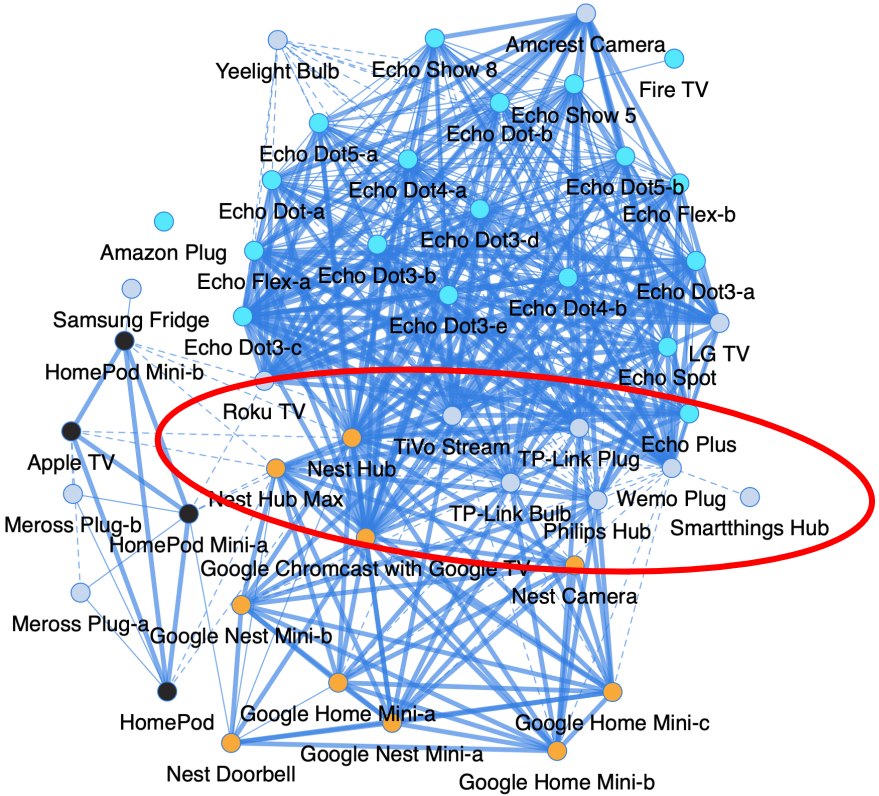

Intra-vendor communication across devices in Amazon, Google, and Apple's ecosystem.

Inter-vendor communication across devices offering interoperable features (e.g., casting, using open-source protocols)

How do these devices interact with each other?

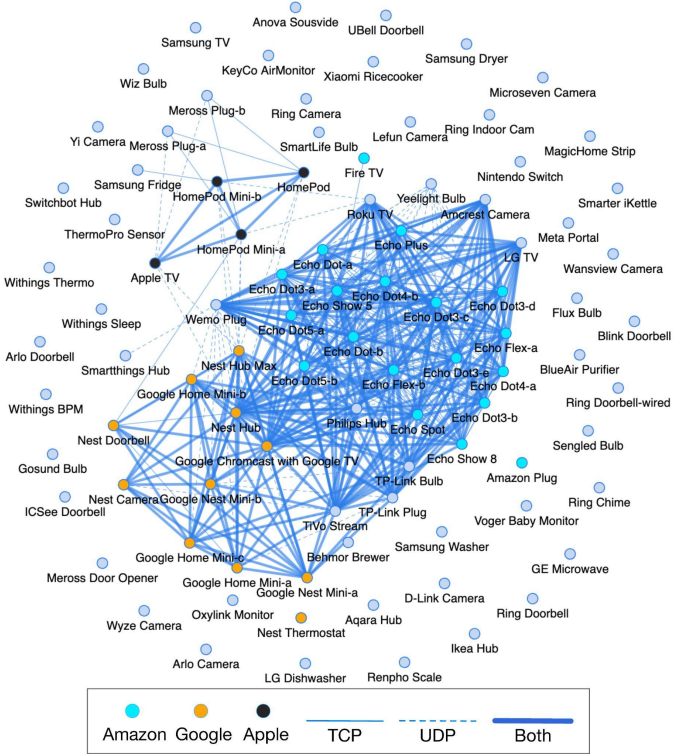

35 different protocols

Nearly half (43/93) devices communicate via unicast

How do these devices interact with each other?

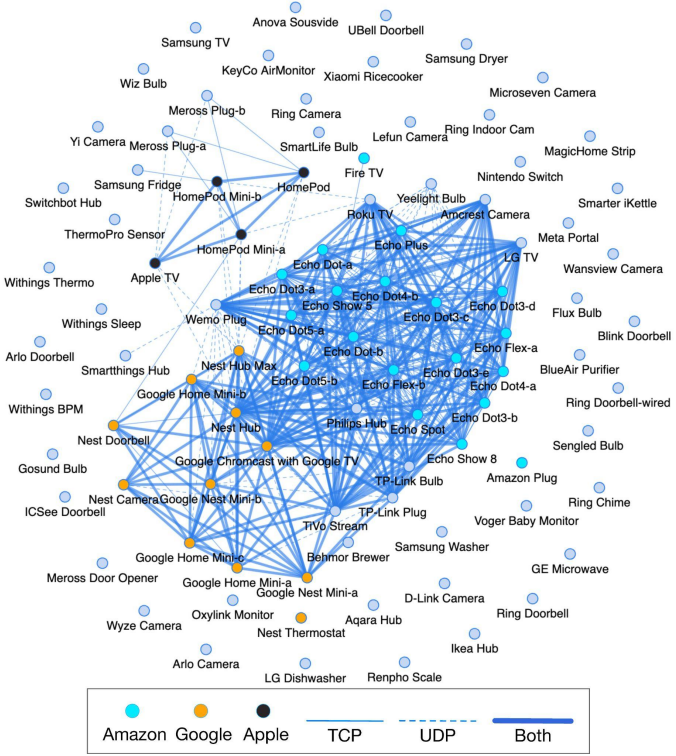

35 different protocols

Nearly half (43/93) devices communicate via unicast

(mostly) Discovery protocols

93% of devices use broadcast-based protocols e.g., ARP, XID/LLC, DHCP.

73% of devices use multicast ones e.g., mDNS, ICMPv6, SSDP, DHCPv6, IGMPv2/v3, CoAP.

What are the privacy and security threats?

Dissemination of sensitive device and network information through discovery protocols

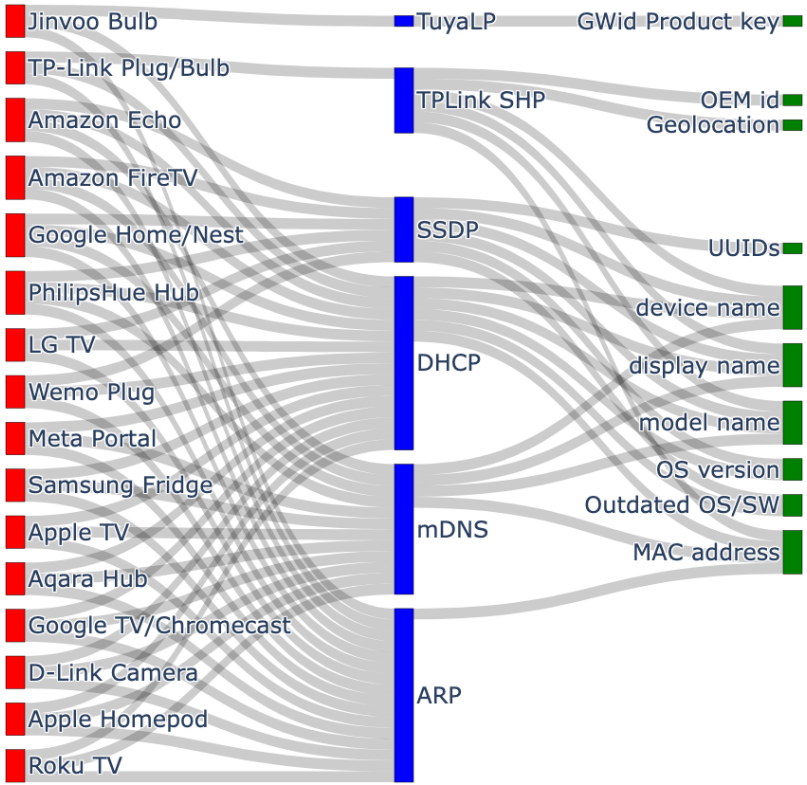

What are the privacy and security threats?

Dissemination of sensitive device and network information through discovery protocols

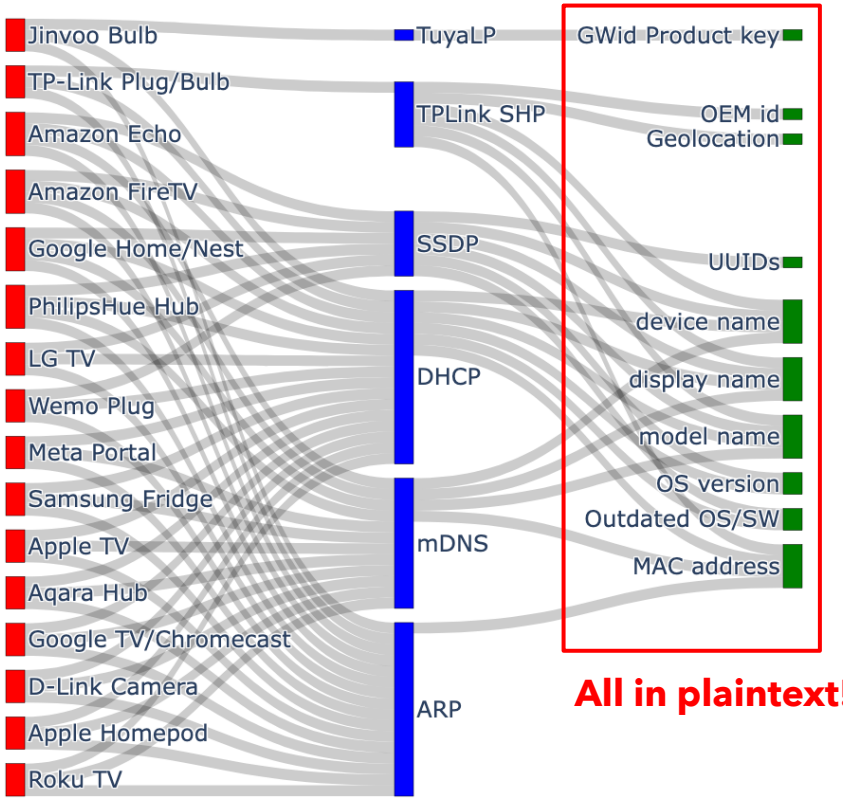

What are the privacy and security threats?

Dissemination of sensitive device and network information through discovery protocols

Check out our paper for more details about other characteristics and security & privacy issues we found.

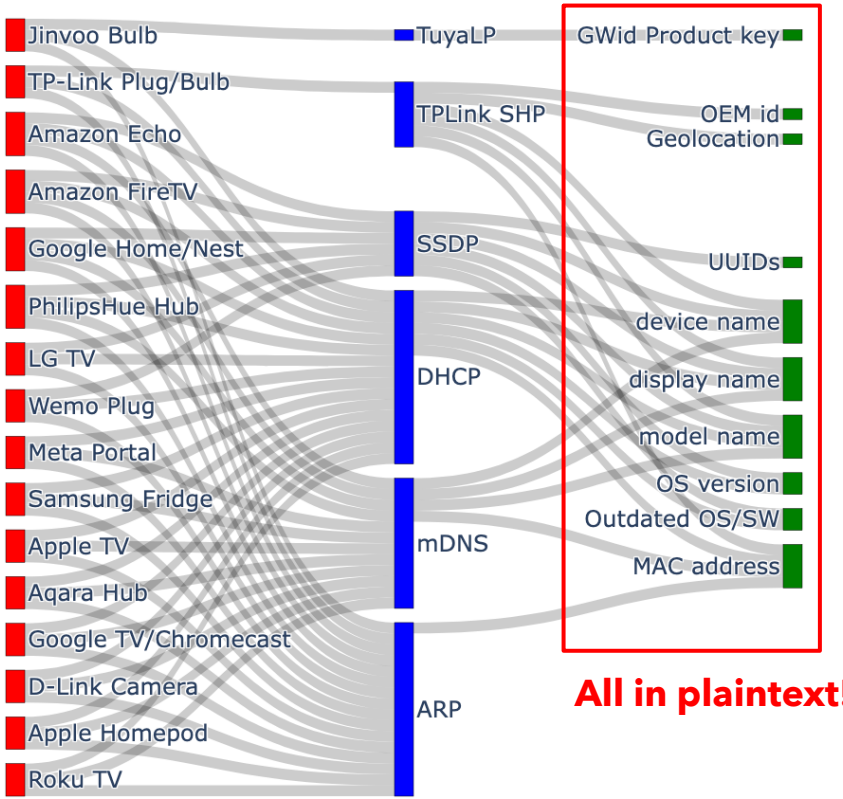

Do advertising and tracking services collect network and device information in the Android platform?

Do advertising and tracking services collect network and device information in the Android platform?

Android

Apps and SDKs can scan the local network and collect information exposed by smart devices using only the INTERNET permission (automatically granted at install time).

No user consent required.

Do advertising and tracking services collect network and device information in the Android platform?

Android

Apps and SDKs can scan the local network and collect information exposed by smart devices using only the INTERNET permission (automatically granted at install time).

No user consent required.

Bypass runtime permission to access WiFi SSID/BSSID:

- Android 13 permission: NEARBY_WIFI_DEVICES
- Pre-Android 13: ACCESS_COARSE_LOCATION or ACCESS_FINE_LOCATION from Android 9

Side-channel

Apps and SDKs harvest local network information

IoT devices relay sensitive information from other devices in local network to mobile apps

Apps and SDKs harvest local network information

IoT devices relay sensitive information from other devices in local network to mobile apps

Apps and SDKs harvest local network information

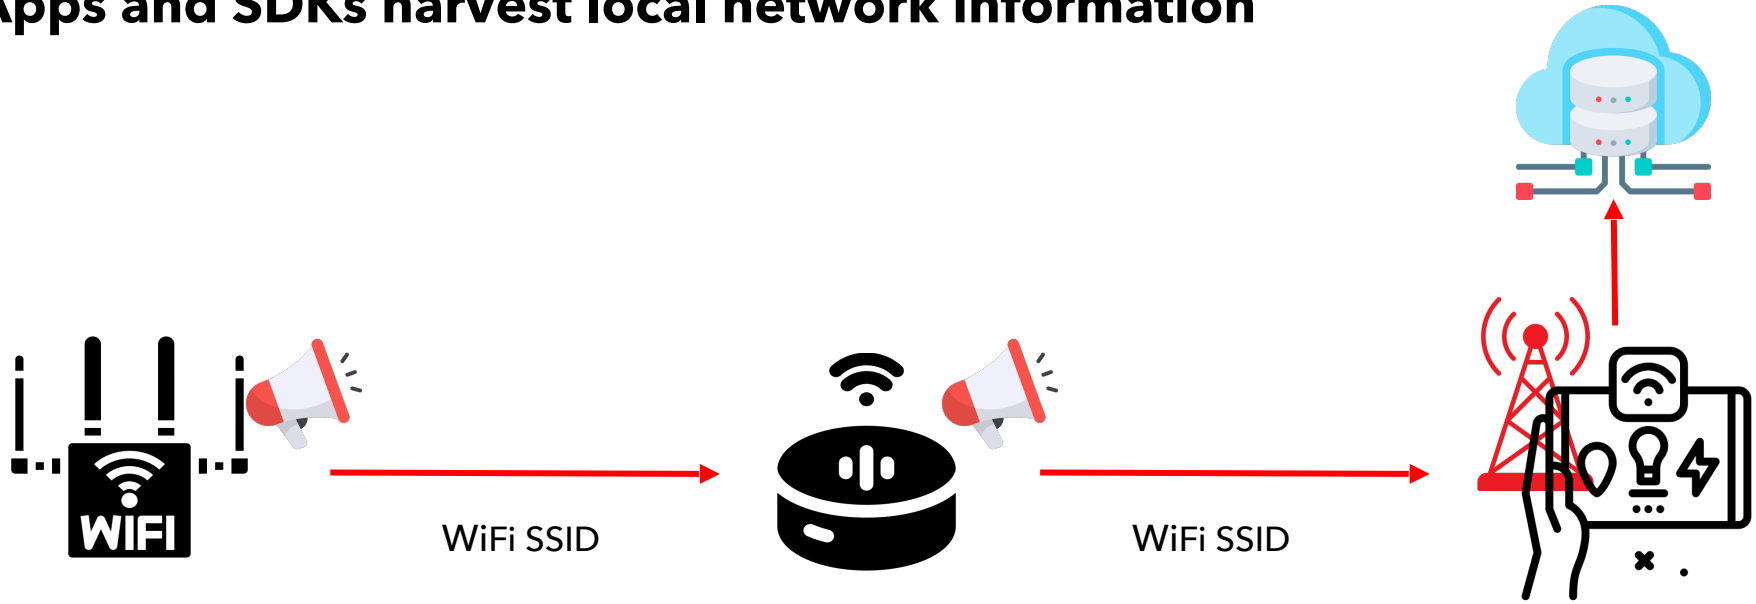

IoT devices relay sensitive information from other devices in local network to mobile apps

Apps and SDKs harvest local network information for advertising & tracking purposes

- **AppDynamics analytics and profiling SDK** collect device information in SSDP/UPnP messages.

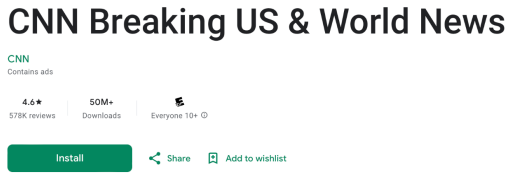


```
HTTP/1.1 200 OK
SERVER: Linux, UPnP/1.0, Private UPnP SDK
...
<?xml version="1.0" ?>
<friendlyName>AMC020SC43PJ749D66</friendlyName>
<serialNumber>9c:8e:cd:0a:33:1b</serialNumber>
<UDN>uuid:device_3_0-AMC020SC43PJ749D66</UDN>
<serviceList>
<service>
```

Apps and SDKs harvest local network information

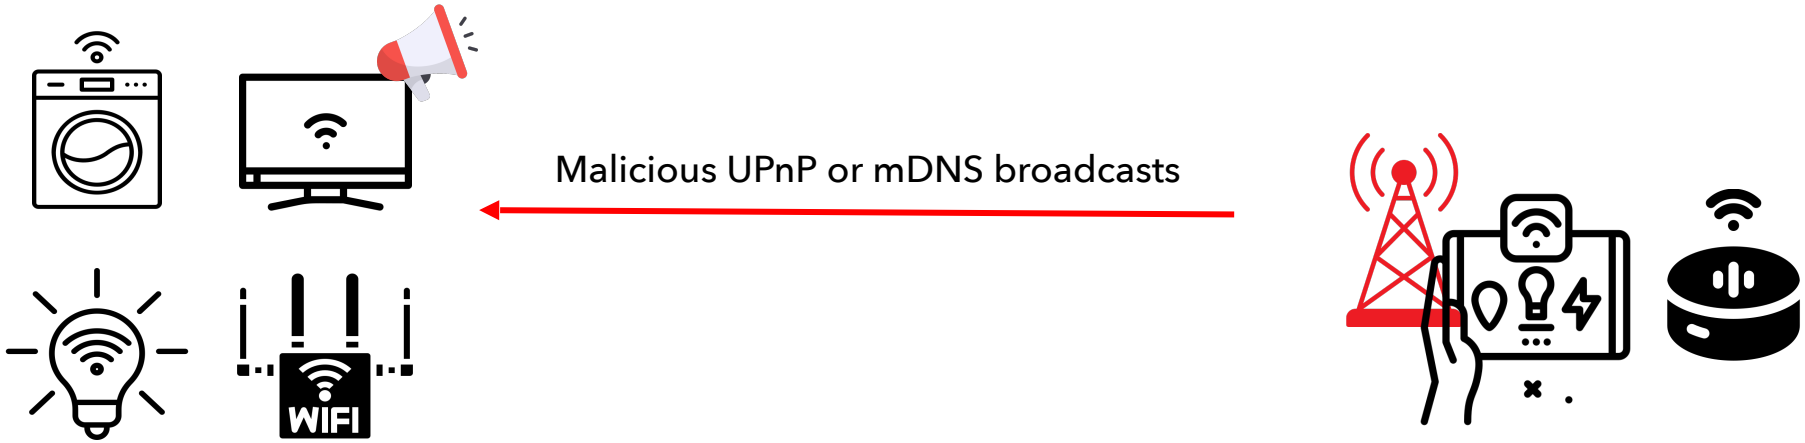

IoT and regular apps & SDK scan and collect MAC address, and WiFi SSID

Apps and SDKs harvest local network information

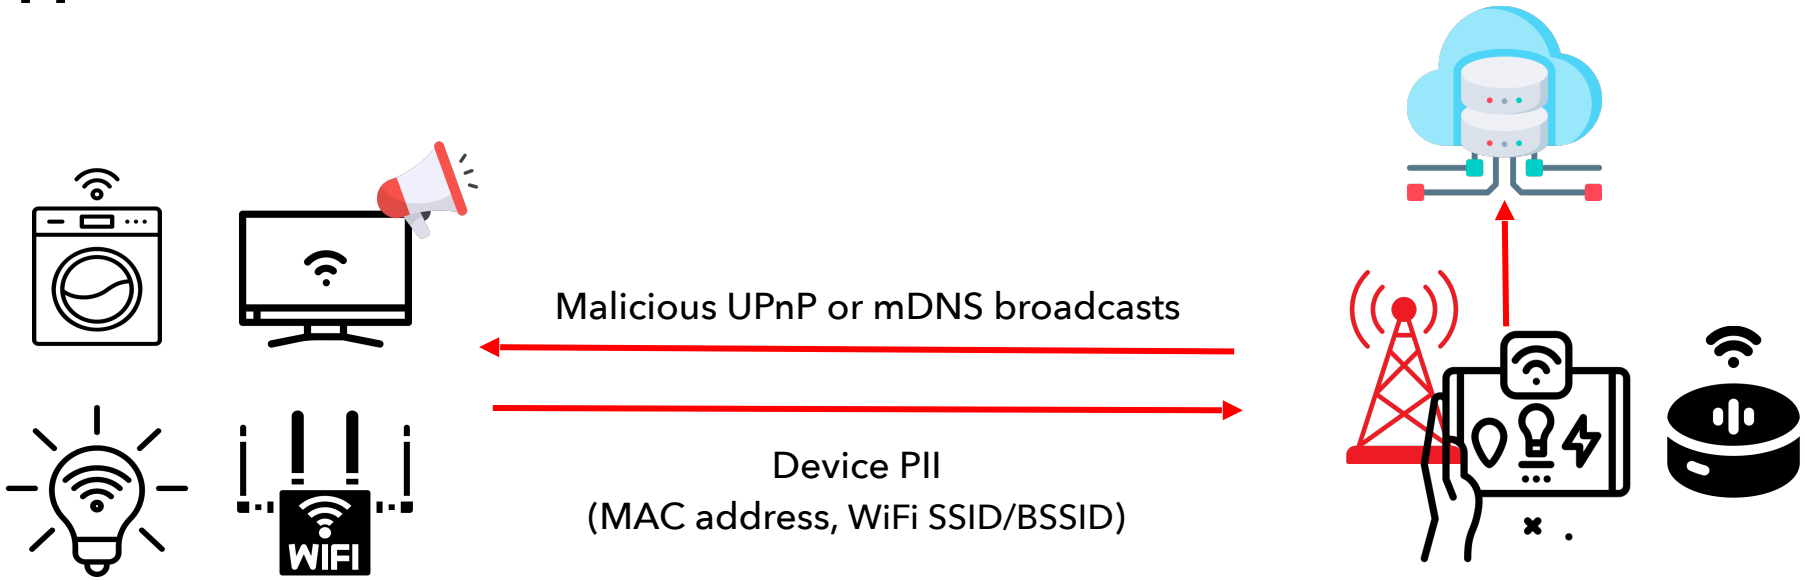

IoT and regular apps & SDK scan and collect MAC address, and WiFi SSID

Apps and SDKs harvest local network information for advertising & tracking purposes

- **Umlaut InsightCore monetization SDK** collects the list of SSDP/UPnP connected devices.

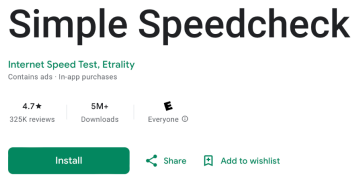


```
    const-string v3, "M-SEARCH * HTTP/1.1\r\nHost: 239.255.255.250:1900
\"ssdp:discover\"\r\nMX: 1\r\nST: urn:schemas-upnp-
org:device:InternetGatewayDevice:1\r\n"

    invoke-virtual {v3}, Ljava/lang/String;→getBytes()[B
    new-instance v5, Ljava/net/DatagramPacket;
    const-string v7, "239.255.255.250"

    invoke-static {v7}, Ljava/net/InetAddress;-
>getByName(Ljava/lang/String;)Ljava/net/InetAddress;
```

Apps and SDKs harvest local network information for advertising & tracking purposes

NetBIOS

- **Innosdk, a third-party anti-cheat and advertising library**

  It sends NetBIOS requests to every IP in the *192.168.0.0/24* prefix and sends local network info to *gw.innotechworld.com* endpoint.

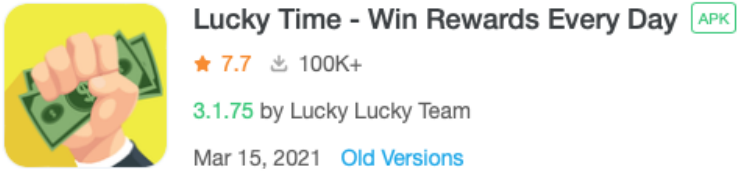

Apps and SDKs harvest local network information for advertising & tracking purposes

NetBIOS

- **Innosdk, a third-party anti-cheat and advertising library**

  It sends NetBIOS requests to every IP in the *192.168.0.0/24* prefix and sends local network info to *gw.innotechworld.com* endpoint.

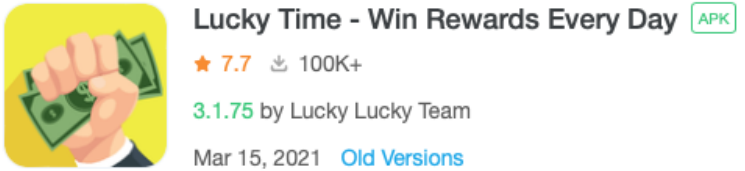

All apps with this SDK have been removed from the Google Play Store

Can exposed local information be used for household fingerprinting and cross-device tracking?

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Higher entropy indicates greater fingerprintability

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Higher entropy indicates greater fingerprintability

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Higher entropy indicates greater fingerprintability

For reference, entropy of HTTP User Agent: **~10.5**

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

# of Identifiers	Entropy
1	6.7
2	14.5
3	20.1

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Higher entropy indicates greater fingerprintability

For reference, entropy of HTTP User Agent: **~10.5**

Exposing all three identifiers makes your household **highly distinctive**

Can exposed local information be used for household fingerprinting and cross-device tracking?

Smart home fingerprintability

IoT Inspector dataset: mDNS and SSDP responses from 12k devices from 3.8k households

3 types of identifiers:
(1) Names, (2) UUIDs, (3) MAC Address

# of Identifiers	Entropy
1	6.7
2	14.5
3	20.1

Metric: entropy to measure web fingerprintability defined by the Electronic Frontier Foundation (EFF)

Higher entropy indicates greater fingerprintability

For reference, entropy of HTTP User Agent: **~10.5**

Exposing all three identifiers makes your household **highly distinctive**

2,814 households exposed UUIDs; 94.2% of these households can be uniquely identified.

Disclosure & Responses from vendors

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.
- We sent reports to 19 IoT vendors regarding potential security issues.

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.
- We sent reports to 19 IoT vendors regarding potential security issues.
- We contacted regulators in relevant jurisdictions regarding potential privacy issues.

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.
- We sent reports to 19 IoT vendors regarding potential security issues.
- We contacted regulators in relevant jurisdictions regarding potential privacy issues.

Google acknowledges this is a real issue and harms users' privacy.
Mitigations: **new permissions** in the Android OS, **app review** processes, and general **IoT standardization** efforts.

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.
- We sent reports to 19 IoT vendors regarding potential security issues.
- We contacted regulators in relevant jurisdictions regarding potential privacy issues.

Google acknowledges this is a real issue and harms users' privacy.
Mitigations: **new permissions** in the Android OS, **app review** processes, and general **IoT standardization** efforts.

11 out of 19 IoT vendors responded to our reports.

Disclosure & Responses from vendors

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.
- We sent reports to 19 IoT vendors regarding potential security issues.
- We contacted regulators in relevant jurisdictions regarding potential privacy issues.

Google acknowledges this is a real issue and harms users' privacy.
Mitigations: **new permissions** in the Android OS, **app review** processes, and general **IoT standardization** efforts.

11 out of 19 IoT vendors responded to our reports.

Signify/Hue: new identifier selected at random to replace the current UUID.

This attack vector is also exploitable by other in-network adversaries

Potential in-LAN adversaries:

- IoT devices (IoT manufacturers, and providers)
- Routers, network service providers
- Smart TV apps
- Visitors, roommates, AirBnB users
- Compromised devices
- …

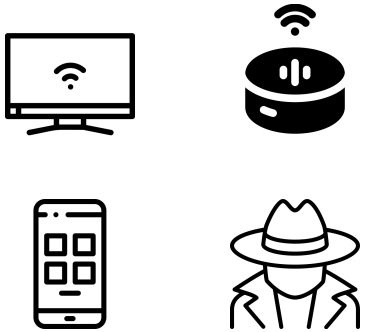

This attack vector is also exploitable by other in-network adversaries

Potential in-LAN adversaries:

- IoT devices (IoT manufacturers, and providers)
- Routers, network service providers
- Smart TV apps
- Visitors, roommates, AirBnB users
- Compromised devices
- …

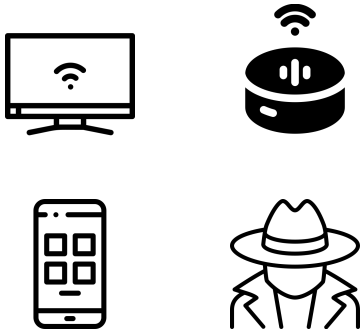

iOS

Network scanning:

- Developers need **explicit approval from Apple** to access multicast sockets.
- **Permission required**: NSLocalNetworkUsageDescription.

  Requests explicit user consent.

Mitigations and Actions

- Developers require explicit approval from the platform for local network access.
- Users can grant or deny local network access via explicit consent and permissions.
- Usable security & privacy controls.

Mitigations and Actions

- Developers require explicit approval from the platform for local network access.
- Users can grant or deny local network access via explicit consent and permissions.
- Usable security & privacy controls.

- Transparency
- Secure-by-design firmware and timely updates
- Supply chain hardening

Mitigations and Actions

- Developers require explicit approval from the platform for local network access.
- Users can grant or deny local network access via explicit consent and permissions.
- Usable security & privacy controls.

- Transparency
- Secure-by-design firmware and timely updates
- Supply chain hardening

- Standardization efforts
- Regulation and certification

Conclusion

- **First characterization:** *local* communication for 93 smart home IoT devices and mobile apps.
- **Sensitive information dissemination**: found in local traffic, including unique IDs, other PII.
- **Fingerprintability and information harvesting**:
  - we demonstrate households are easily fingerprinted, enabling cross-device tracking.
  - we find mobile apps and third-party SDKs harvesting local network information.
- **Disclosure**: We identified responsible parties, ongoing efforts for remediation.

Thank you!

Aniketh Girish
aniketh.girish@imdea.org

Datasets and code available here: https://github.com/Android-Observatory/IoT-LAN

Backup

Disclosure

- We reported the Android side channel issue to Google.
- We provided a list of misbehaving Android apps to Google.
- We sent reports to 19 IoT vendors regarding potential security issues.
- We contacted regulators in relevant jurisdictions regarding potential privacy issues.

We privately inform responsible parties through their vulnerability disclosure programs or customer contacts

We gave vendors 30 days notice given timing constraints for publication

How these devices interact with each others?

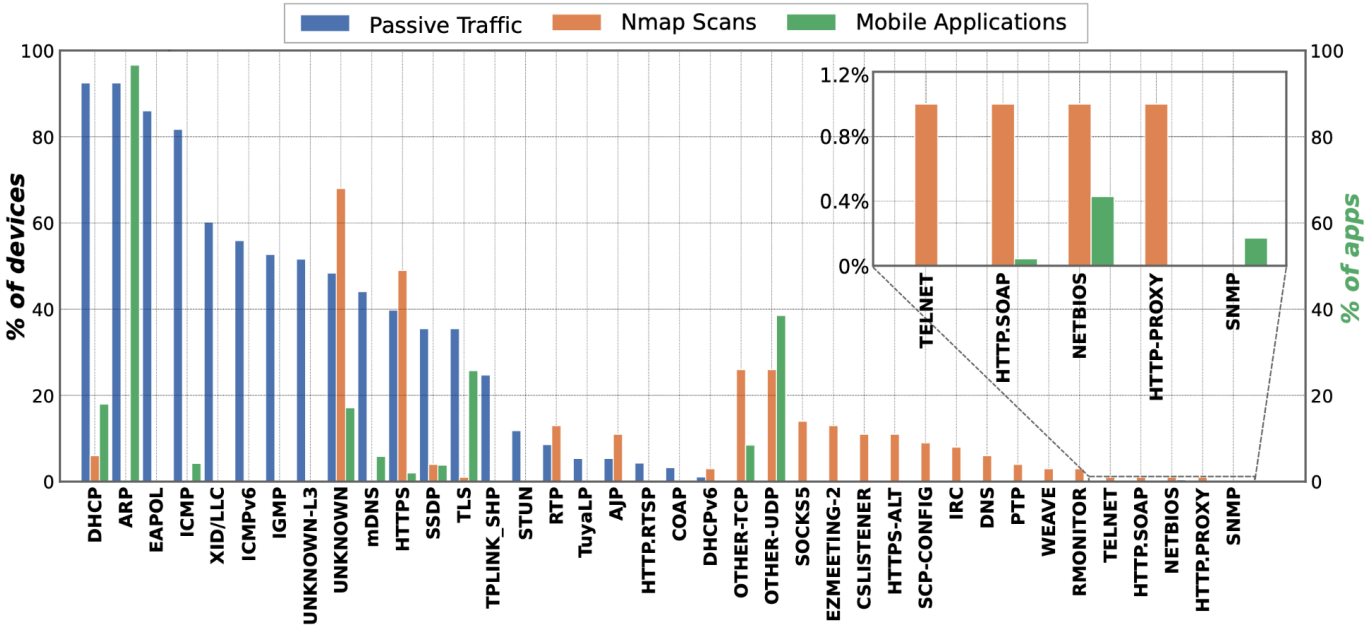

Nearly half (43/93) devices use **TCP or UDP unicast communication**

How these devices interact with each others?

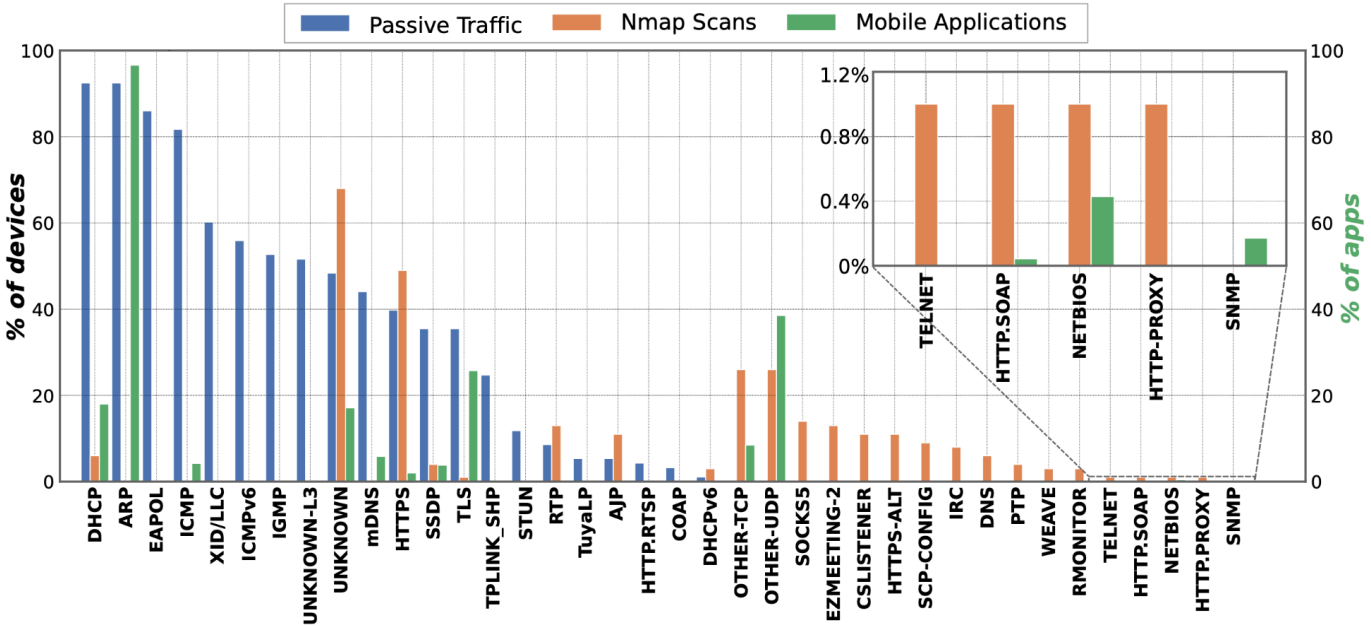

Nearly half (43/93) devices use **TCP or UDP unicast communication**

(mostly) Discovery protocols

- **93% of devices use broadcast**-based protocols like ARP, XID/LLC, DHCP
- **73% of devices use multicast** ones like mDNS, ICMPv6, SSDP, DHCPv6, IGMPv2/v3, and CoAP.

SSDP	HTTP/1.1 200 OK SERVER: Linux, UPnP/1.0, Private UPnP SDK ... <?xml version="1.0" ?> <friendlyName>AMC020SC43PJ749D66</friendlyName> <serialNumber>9c:8e:cd:0a:33:1b</serialNumber> <UDN>uuid:device_3_0-AMC020SC43PJ749D66</UDN> <serviceList> <service>
mDNS	Ethernet II, Src: PhilipsL_68:5f:61 (00:17:88:68:5f:61), Dst: IPv4mcast_fb (01:00:5e:00:00:fb) ... Multicast Domain Name System (response) Philips Hue - 685F61._hue._tcp.local: type TXT, class IN, cache flush _hue._tcp.local: type PTR, class IN, Philips Hue - 685F61._hue._tcp.local 1.6.F.5.8.6.E.F.F.F.8.8.7.1.2.0.0.0.0.0.0.0.0.0.0.0.0.0.8.E.F.ip6.arpa: type PTR
TPLINK-SHP	{"system":{"get_sysinfo ... "deviceId":"8006E8E9017F556D283C850B4E29BC1F185334E5", "hwId":"60FF6B258734EA6880E186F8C96DDC61" ... oemId":"FFF22CFF774A0B89F7624BFC6F50D5DE "alias":"TP-Link Plug","dev_name":"Wi-Fi Smart Plug With Energy Monitoring" ... "latitude":42.337681,"longitude":-71.087036
Co-located devices leaking data to the cloud	HTTP/1.1 200 OK {"entity":{"entityId":"SKILL_eyJza2lsbElkIjoiYW16bjEuYXNrLnNraWxsLmI0YmYyYjRkLT -> 8012A5191D2CB6983983DB807412997E18990EFF> –> Light bulb deviceId ","entityType":"CLOUD_DISCOVERED_DEVICE"},"capabilityStates": ["{\"namespace\":\"Alexa.BrightnessController\",\"name\":\"brightness\",\"value\":100,

What are the security and privacy threats?

What are the risks of these information exposure?

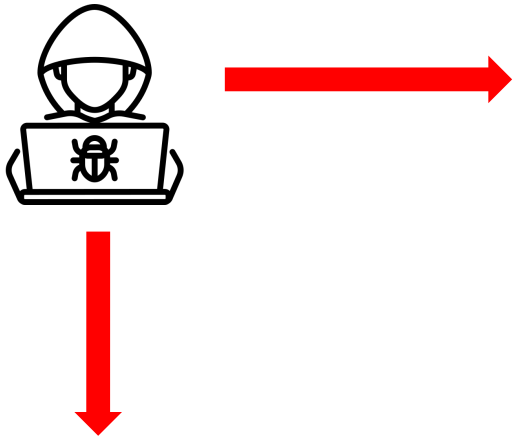

What are the security and privacy threats?

What are the risks of these information exposure?

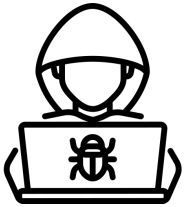

Cross-device tracking & Household and user profiling using

- MAC address
- SSID
- Device model and name
- Services supported, e.g., printing
- UUIDs
- Geolocation
- Device display name, e.g., Peter's Apple TV
- …

Targeted attacks using

- Device model
- Software component version
- OS version
- UUIDs
- Services supported, e.g., printing
- …

What are the security and privacy threats?

What are the risks of these information exposure?

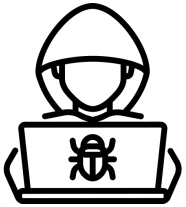

Cross-device tracking & Household and user profiling using

- MAC address
- SSID
- Device model and name
- Services supported, e.g., printing
- UUIDs
- Geolocation
- Device display name, e.g., Peter's Apple TV
- …

infer

- **Household social structures and socio-economic level such as your household type, income level, parantship/ relationship status, etc.**
- **Geolocation of the household**

Targeted attacks using

- Device model
- Software component version
- OS version
- UUIDs
- Services supported, e.g., printing
- …

Game Console	Generic IoT	Home Appliance	Home Automation	Media/TV	Surveillance	Voice Assistant
Nintendo (1)	Keyco (1)	Anova (1)	Amazon (1)	Amazon (1)	Amcrest (1)	Amazon (17)
	Oxylink (1)	Behmor (1)	Aqara (1)	Apple (1)	Arlo (2)	Apple (3)
	Renpho (1)	Blueair (1)	Google (1)	Google (1)	Blink (1)	Meta (1)
	Tuya (1)	GE (1)	IKEA (1)	LG (1)	D-Link (1)	Google (7)
	Withings (3)	LG (1)	MagicHome (1)	Roku (1)	Google (2)	
		Samsung (3)	Meross (3)	Samsung (1)	ICSee (1)	
		Smarter (1)	Philips (1)	Tivostream (1)	Lefun (1)	
		Xiaomi (1)	Ring (1)		Microseven (1)	
			Sengled (1)		Ring (4)	
			SmartThings (1)		Tuya (1)	
			SwitchBot (1)		Ubell (1)	
			TP-Link (2)		Wansview (1)	
			Tuya (3)		Wyze (1)	
			WeMo (1)		Yi (1)	
			Wiz (1)			
			Yeelight (1)			

Table 3: IoT devices under test categorized by device type. The number in the parentheses indicates the number of devices.

#	Pdt	Vdr	Dev	Σ Hse	Identifier(s)	Hse	Ent
0	154	107	4,175	1,811	N/A	N/A	N/A
1	160	100	6,915	3,007	name	2 (50.0%)	3.4
					UUID	2,814 (94.2%)	8.9
					MAC	572 (94.4%)	7.8
2	76	59	1,577	1,201	name, UUID	22 (81.8%)	12.3
					UUID, MAC	1,182 (95.6%)	16.7
3	1	1	2	2	name, UUID, MAC	2 (100.0%)	20.1

Information exposed via mDNS and SSDP.
counts identifier types exposed, including first names, UUIDs, and MAC addresses. **Pdt** counts distinct products exposing this information.
Vdr counts vendors across these products
Dev counts distinct devices
Hse' counts households for these devices.
Identifier(s) column shows which identifier(s) are exposed over how many

Backup end