Lecture #2: IoT and Internet Core Protocols

Cristian Hesselman, Elmer Lastdrager,
Ramin Yazdani, Etienne Khan, and Ting-Han Chen

University of Twente | May 3, 2023

UNIVERSITY OF TWENTE.

SIDN LABS

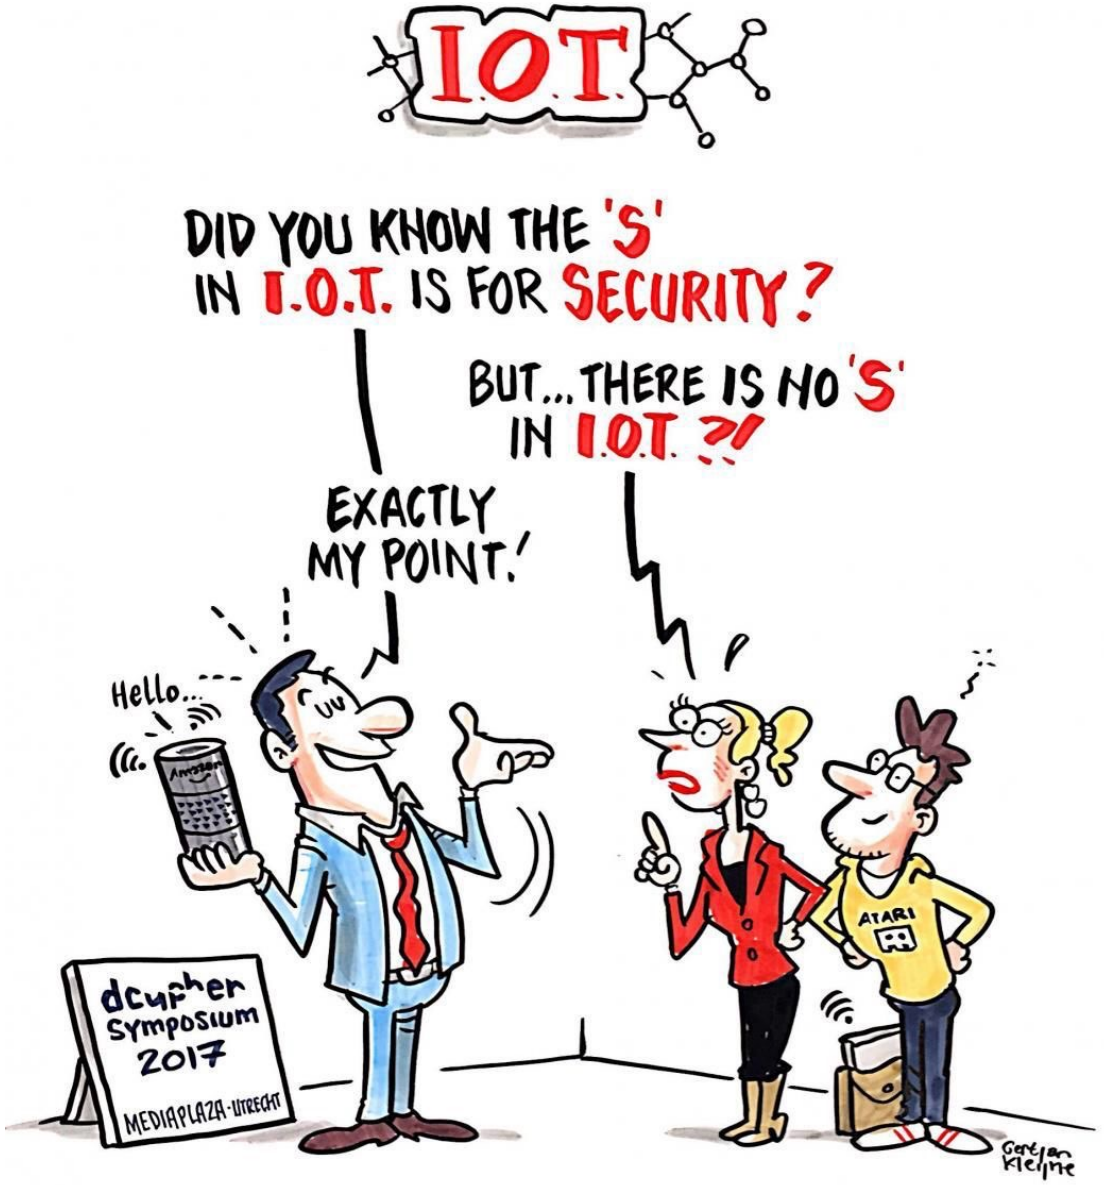

UNIVERSITY OF TWENTE.

SDN LABS

Today's agenda

- Admin
- Introduction to today's lecture
- Paper on the DNS in IoT
- Paper on IPv6 port scanning
- Feedback

Admin

Interactive lectures

- Overall objective: enable you to learn from each other and further increase your understanding of the papers, contributes to preparing yourself for the oral exam

- Interactive format
  - Teachers summarize two papers per lecture
  - Multiple-choice and open questions (not graded) and discussion
  - Enables you to learn from each other, so mandatory to participate

- A 7th “re-sit” lecture in case you miss a lecture (optional for everybody else), same format

Paper summaries

- You must have handed in your two summaries **before 7AM on the day of the lecture**
- Each summary can be at most 250 words, at most 1 single-sided A4 page
- You can add figures, and graphs from the paper or add your own if you like (e.g., concept maps)
- You can use the summaries during the oral exam
- Submit through CANVAS
- You **cannot** complete SSI without submitting 12 paper summaries!

Schedule

No.	Date	Contents
1	Apr 26	Course introduction
2	May 3	Lecture: IoT and Internet Core Protocols
3	May 10	Lecture: IoT Botnet Measurements 1
4	May 17	Lecture: IoT Edge Security Systems
5	???	Guest lecture #1: TBD
6	May 24	Lecture: IoT Device Security
7	May 31	Lecture: IoT Botnet Measurements 2
8	Jun 7	Lecture: IoT in Non-Carpeted Areas
9	???	Guest lecture #2: TBD
10	Jun 14	Lecture: IoT Honeypots (re-sit)

Important dates

- Two summaries per lecture: before the lecture (07:00 CEST) in which the papers will be discussed
- Lab report (PDF) and required files: **Friday June 23, 2023, 23:59 CEST**
- All to be submitted through CANVAS

Introduction to today's lecture

Motivation: impact of insecure IoT devices

[Mirai] [SPIN]

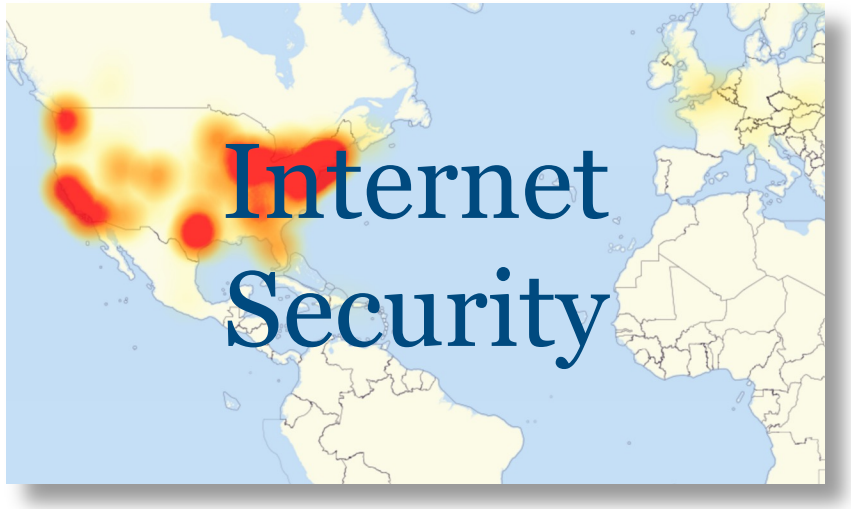

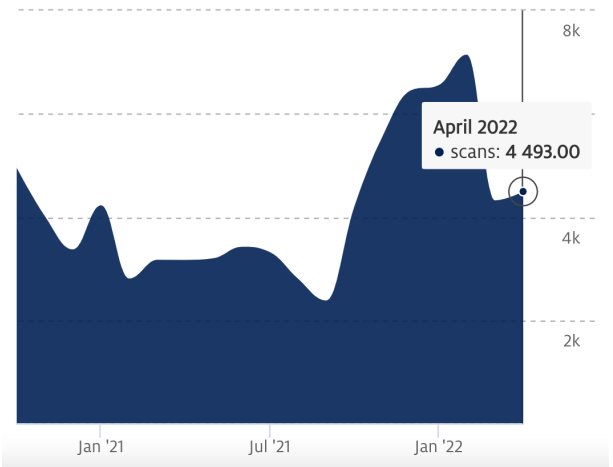

https://stats.sidnlabs.nl/en/security.html#mirai%20scans

[Castle] [SPIN]

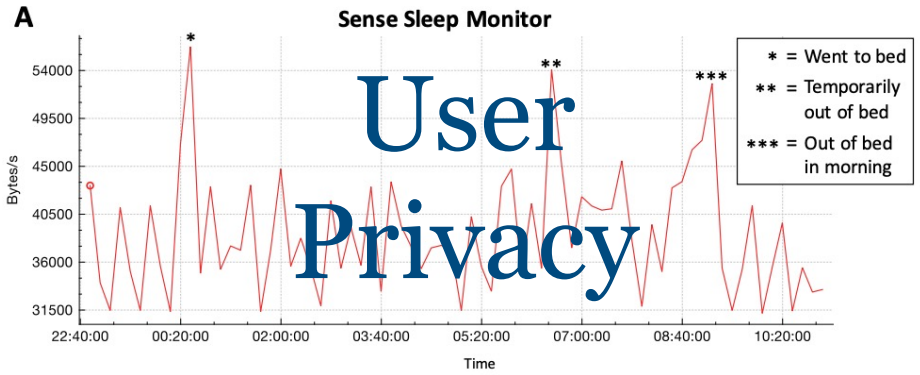

User Privacy

[DNSIoT] [Lora] [Traffic] [2stic.nl]

User Safety

UNIVERSITY OF TWENTE.

SIDN LABS

Today's papers

[DNSIoT] C. Hesselman, M. Kaeo, L. Chapin, kc claffy, M. Seiden, D. McPherson, D. Piscitello, A. McConachie, T. April, J. Latour, and R. Rasmussen, "The DNS in IoT: Opportunities, Risks, and Challenges", IEEE Internet Computing, Vol. 24, No. 4, July-Aug 2020

[IPv6] P. Richter, O. Gasser, and A. Berger, "Illuminating large-scale IPv6 scanning in the internet", In Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22), New York, NY, USA, 410–418, 2022, https://doi.org/10.1145/3517745.3561452.

Today's learning objective

- After the lecture, you will be able to discuss the role of DNS for IoT and challenges of IPv6 scanning

- Not very technical, but important to "set the scene" for more technical papers later in the course (we'll point you to them)

- Contributes to SSI learning goal #1: "Understand IoT concepts and applications, security threats, technical solutions, and a few relevant standardization efforts in the IETF"

“The DNS in IoT: Opportunities, Risks, and Challenges”

IEEE Internet Computing, July-Aug 2020

UNIVERSITY OF TWENTE.

SIDN LABS

Learning Goals

IoT with DNS

Opportunities, Risks, Challenges

UNIVERSITY OF TWENTE.

SIDN LABS

IoT Definition

No Browser. Widely Heterogeneous. Longevity. Background

Let's see what's going on recently

Smart lamp with Emotion

Tablet for IoT control

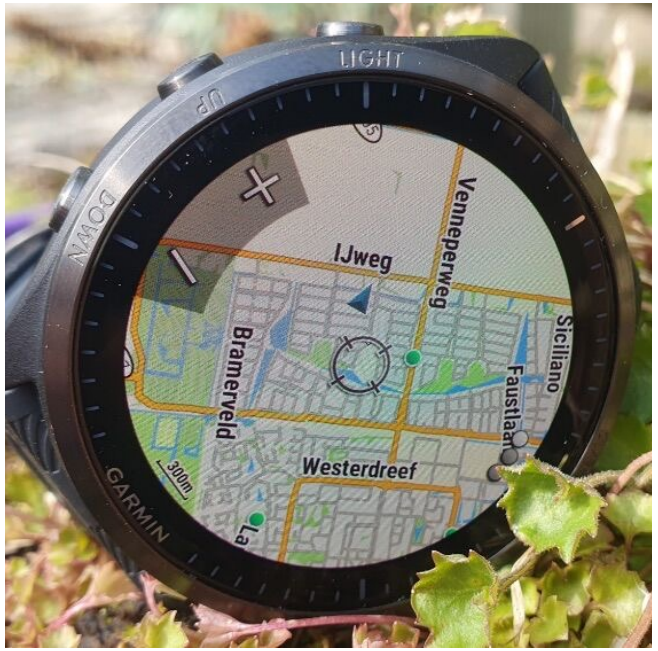

Wristwatch with GPS/LTE

UNIVERSITY OF TWENTE.

SDN LABS

IoT deployments and the Domain Name System (DNS)

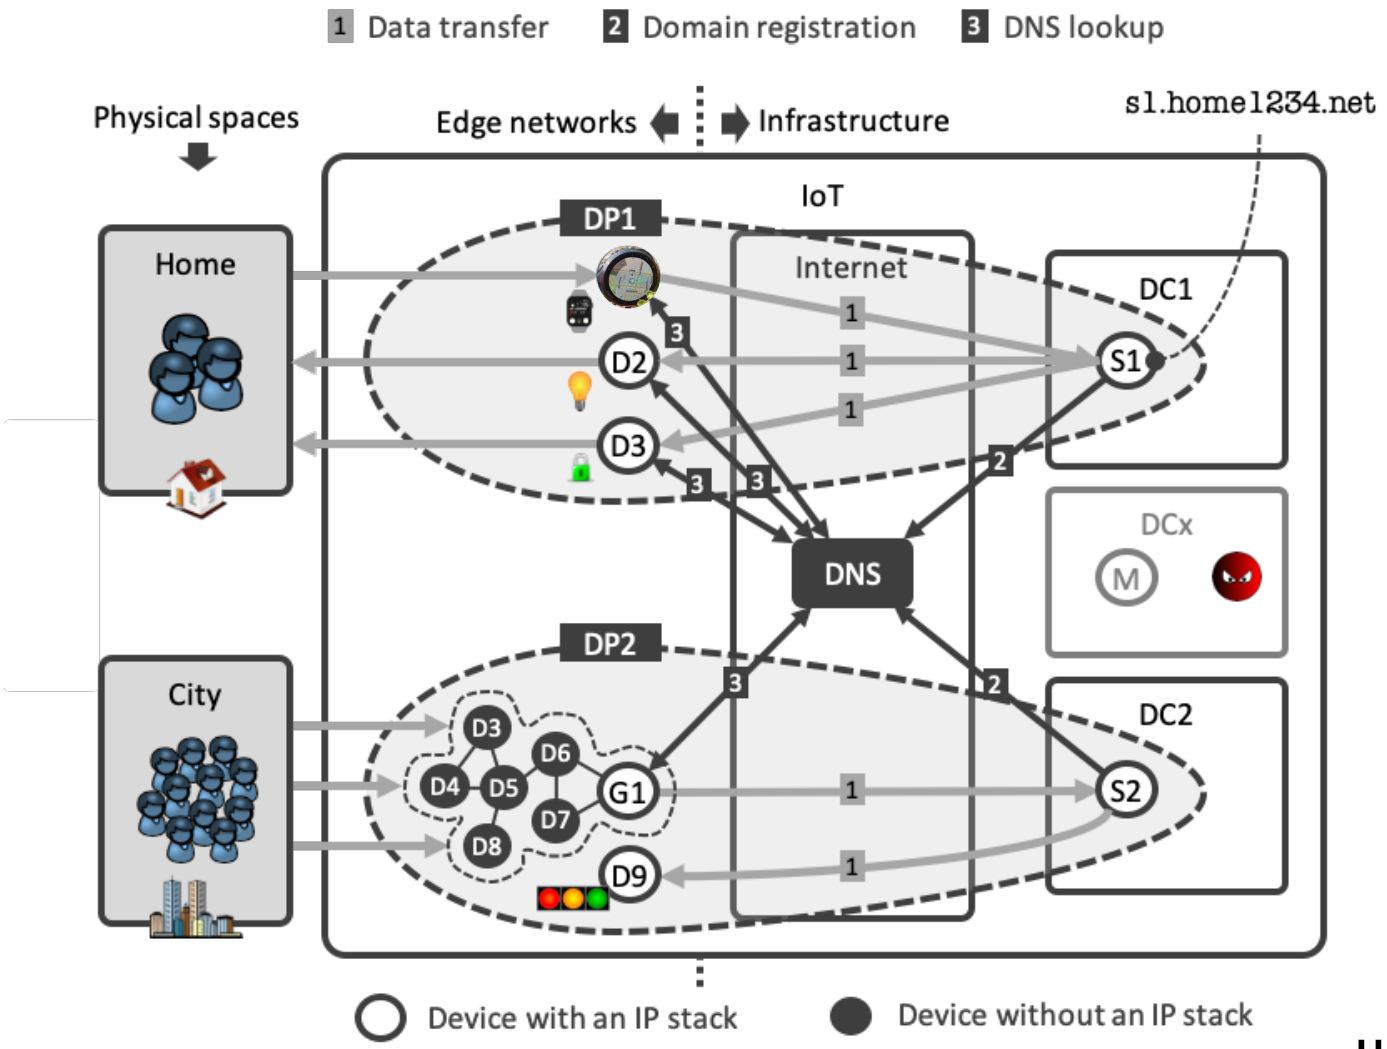

UNIVERSITY OF TWENTE.

SDN LABS

IoT deployments and the Domain Name System (DNS)

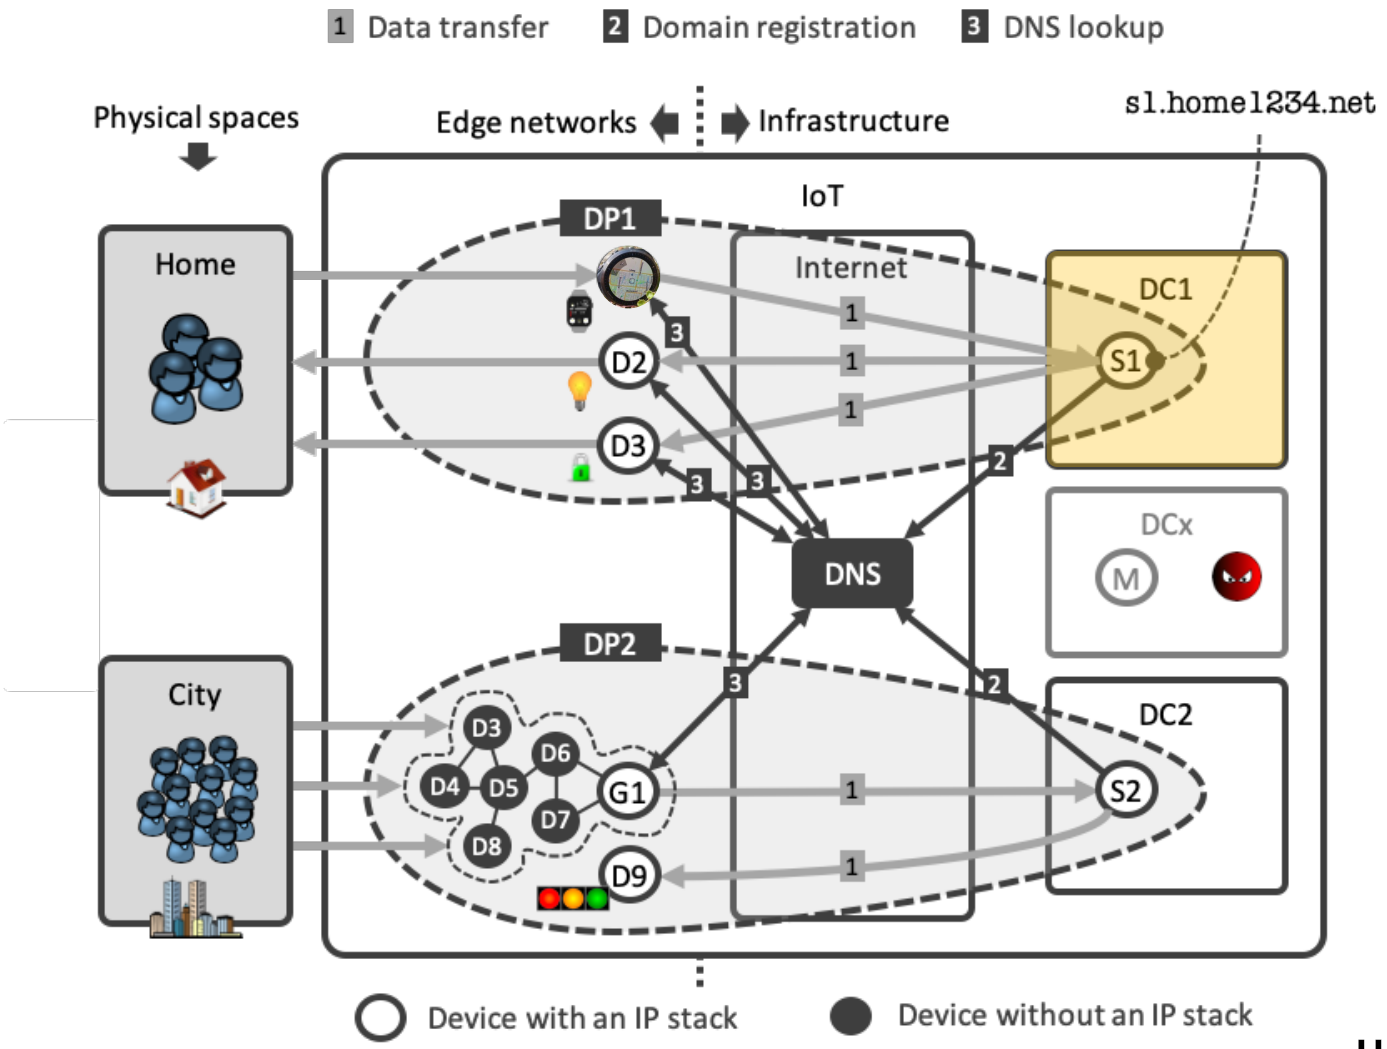

UNIVERSITY OF TWENTE.

SDN LABS

IoT deployments and the Domain Name System (DNS)

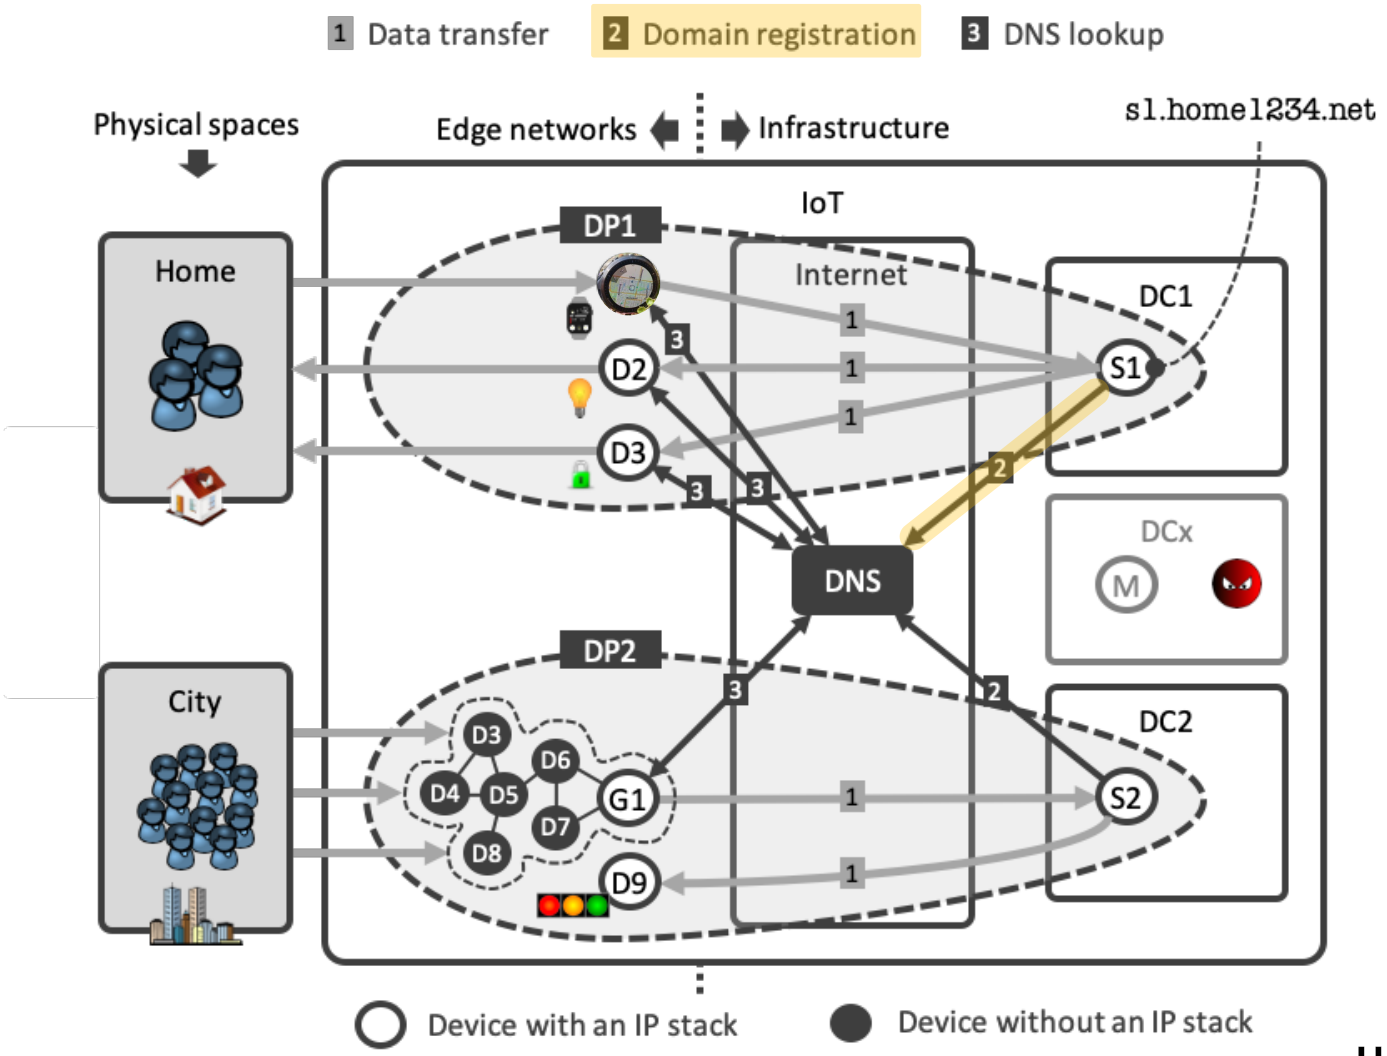

UNIVERSITY OF TWENTE.

SDN LABS

IoT deployments and the Domain Name System (DNS)

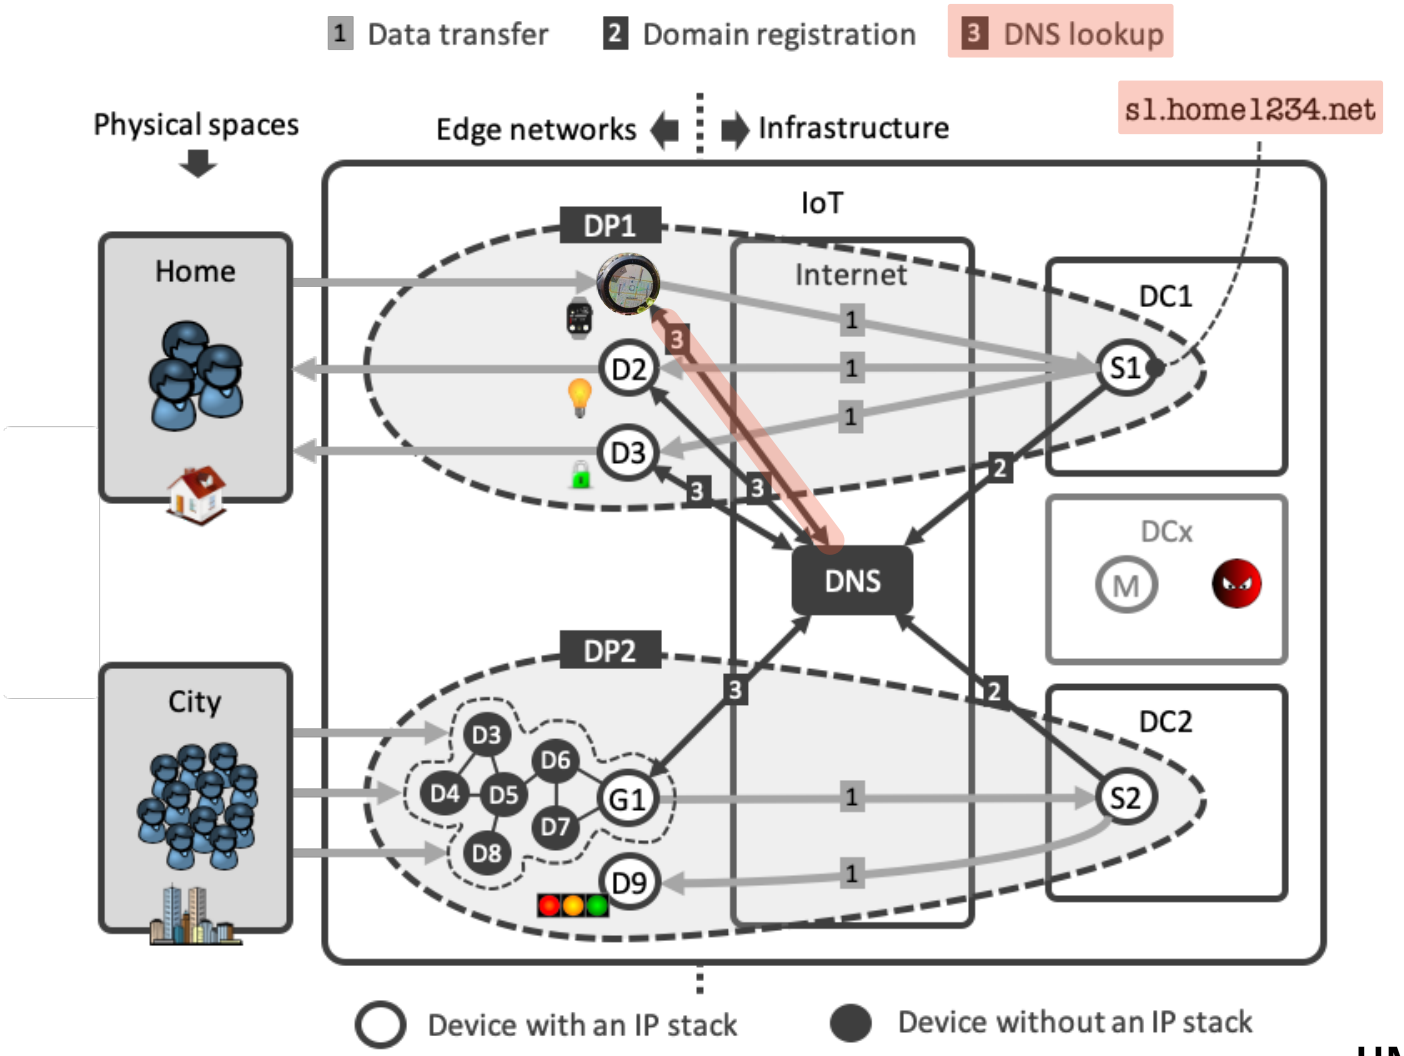

UNIVERSITY OF TWENTE.

SDN LABS

DNS high-level operation

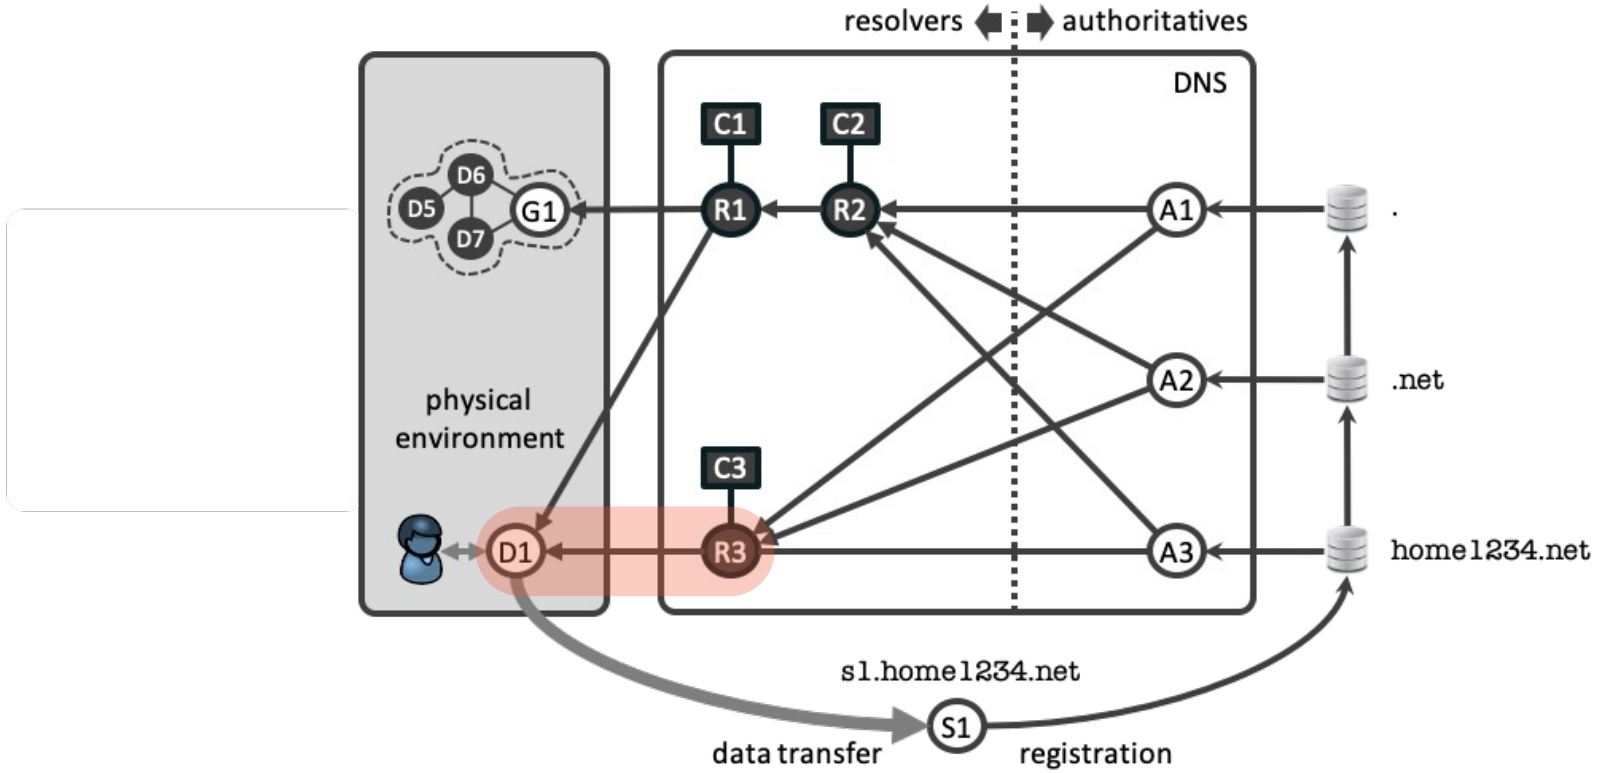

O. van der Toorn, M. Mueller, S. Dickinson, C. Hesselman, A. Sperotto, and R. van Rijswijk-Deij, **"Addressing the Challenges of Modern DNS: A Comprehensive Tutorial"**, Elsevier Computer Science Review, 2022 (to appear)

UNIVERSITY OF TWENTE.

SIDN LABS

DNS high-level operation

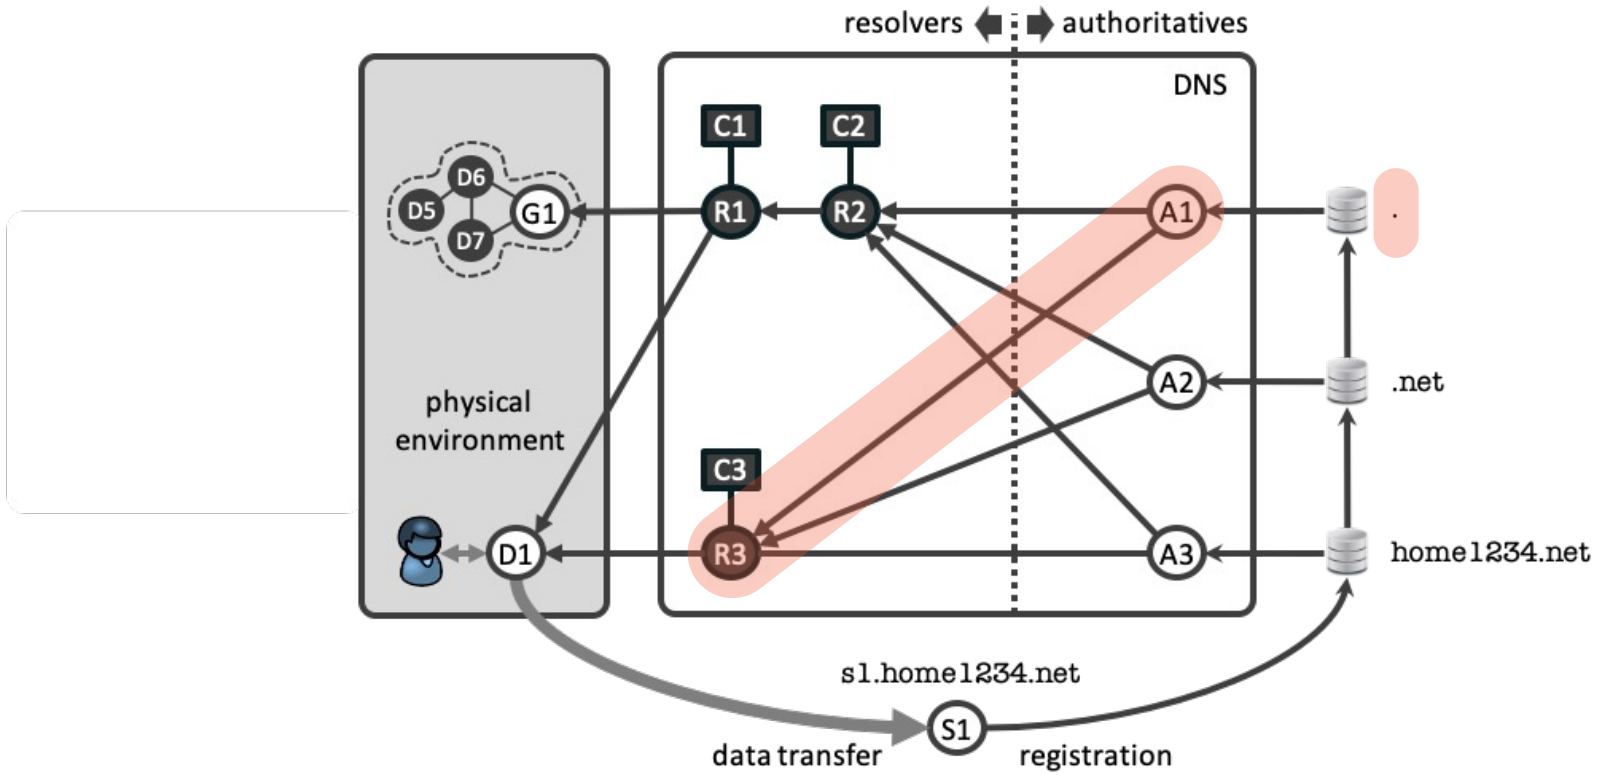

O. van der Toorn, M. Mueller, S. Dickinson, C. Hesselman, A. Sperotto, and R. van Rijswijk-Deij, **"Addressing the Challenges of Modern DNS: A Comprehensive Tutorial"**, Elsevier Computer Science Review, 2022 (to appear)

UNIVERSITY OF TWENTE.

SIDN LABS

DNS high-level operation

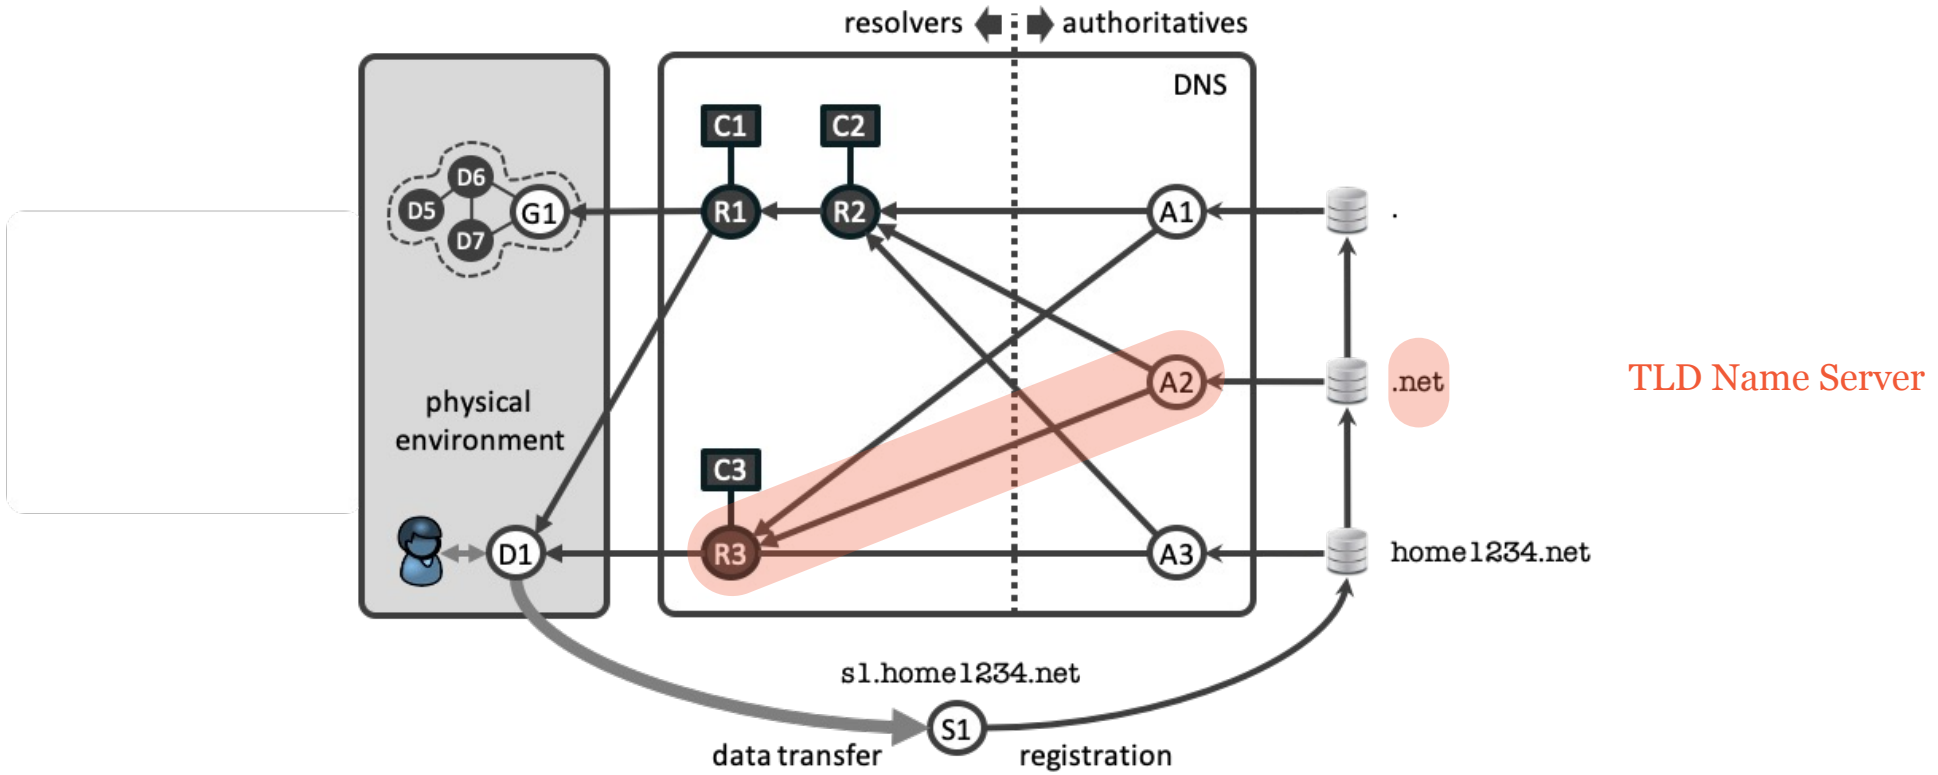

TLD Name Server

O. van der Toorn, M. Mueller, S. Dickinson, C. Hesselman, A. Sperotto, and R. van Rijswijk-Deij, **"Addressing the Challenges of Modern DNS: A Comprehensive Tutorial"**, Elsevier Computer Science Review, 2022 (to appear)

UNIVERSITY OF TWENTE.

SIDN LABS

DNS high-level operation

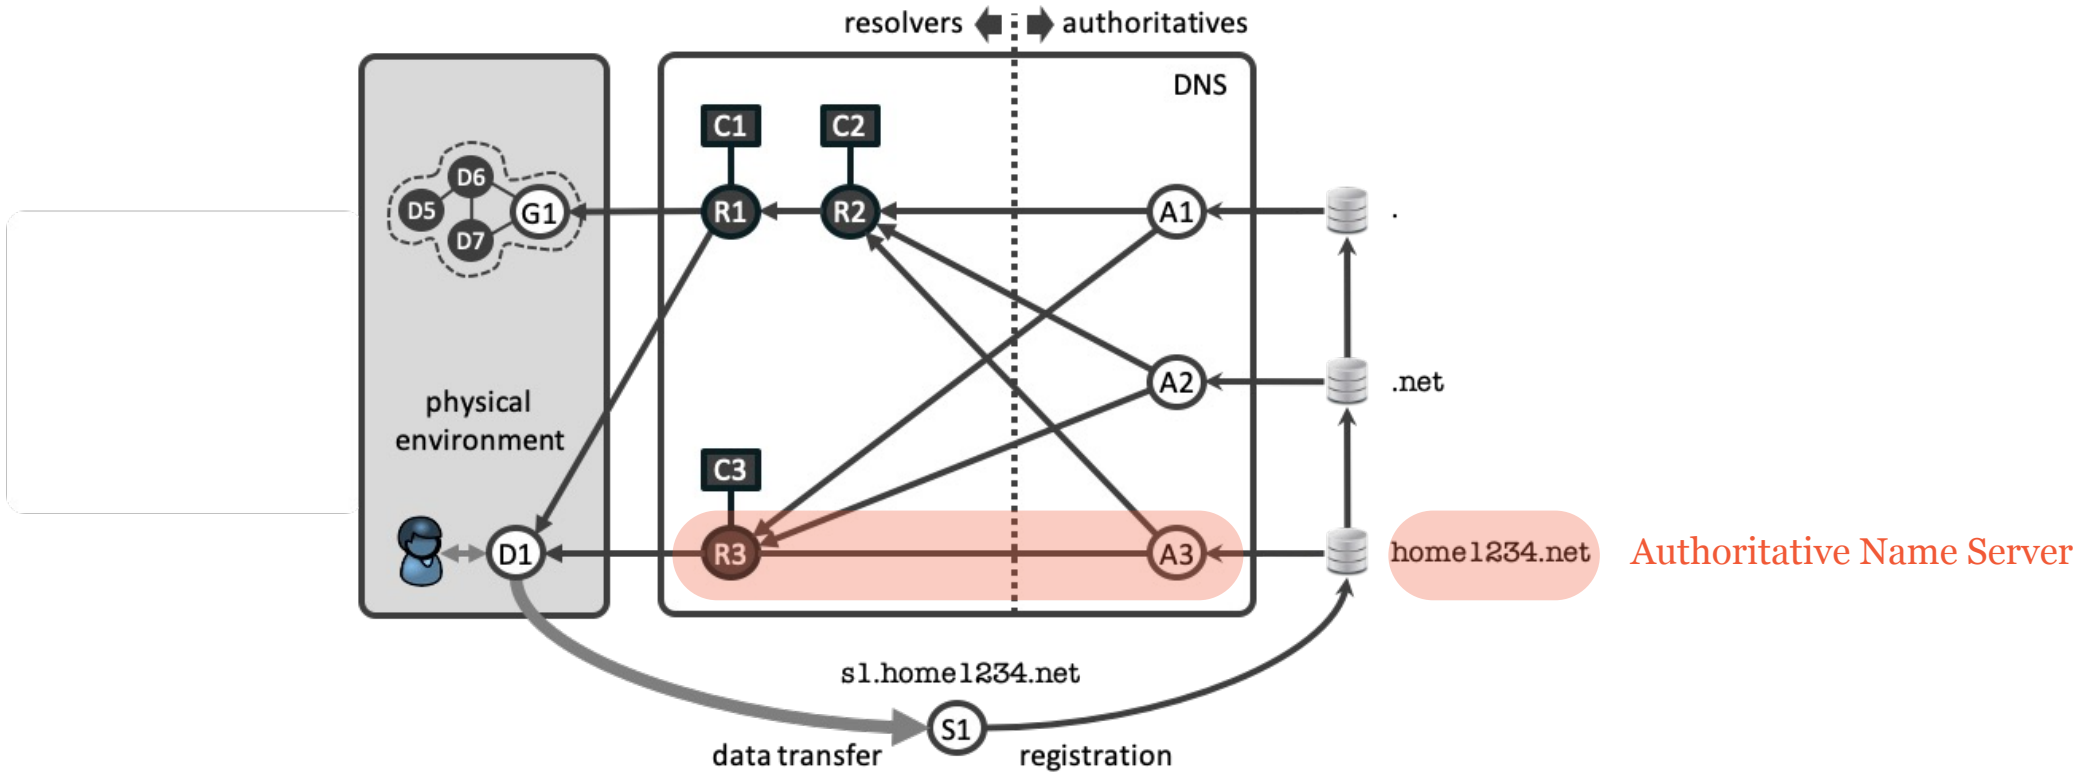

O. van der Toorn, M. Mueller, S. Dickinson, C. Hesselman, A. Sperotto, and R. van Rijswijk-Deij, **"Addressing the Challenges of Modern DNS: A Comprehensive Tutorial"**, Elsevier Computer Science Review, 2022 (to appear)

UNIVERSITY OF TWENTE.

SIDN LABS

IoT deployments and the Domain Name System (DNS)

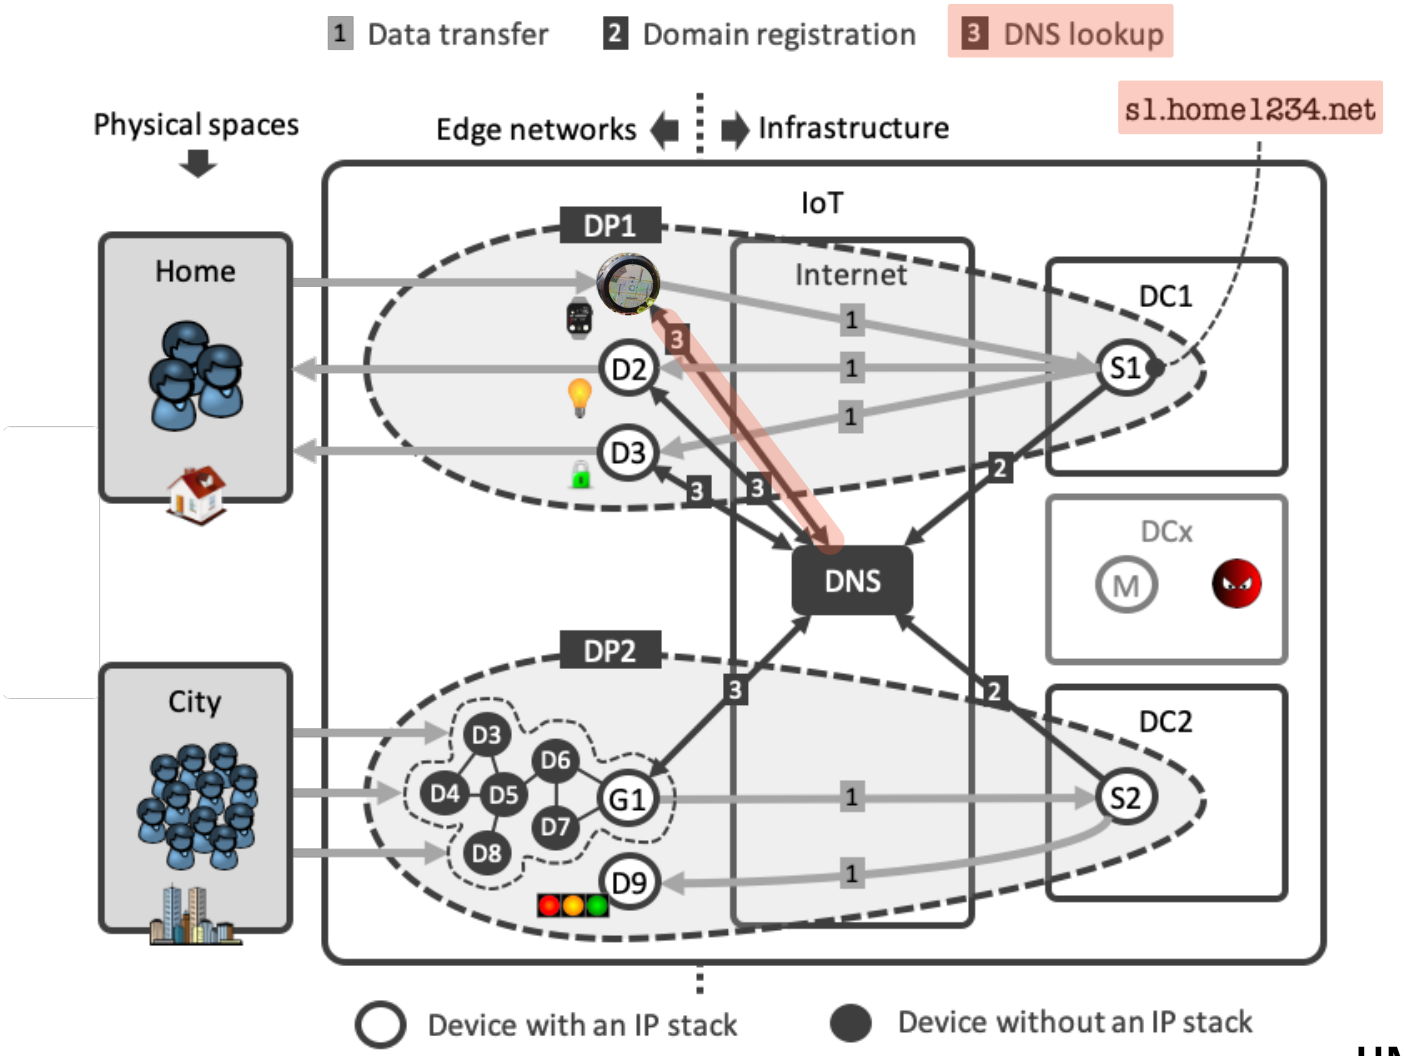

UNIVERSITY OF TWENTE.

SDN LABS

IoT deployments and the Domain Name System (DNS)

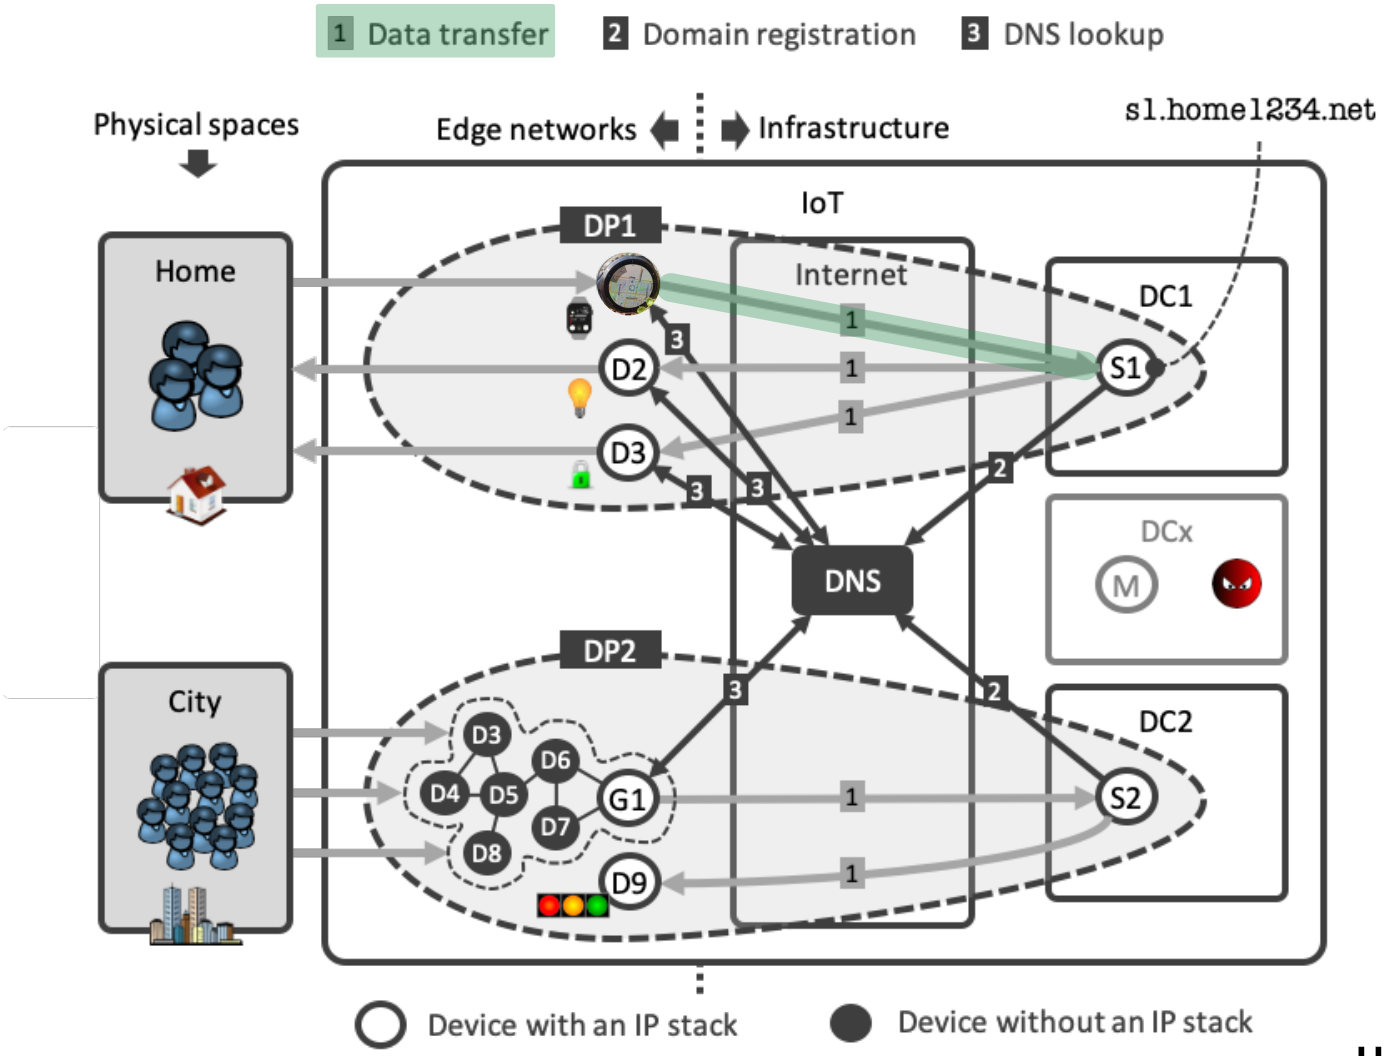

IoT deployments and the Domain Name System (DNS)

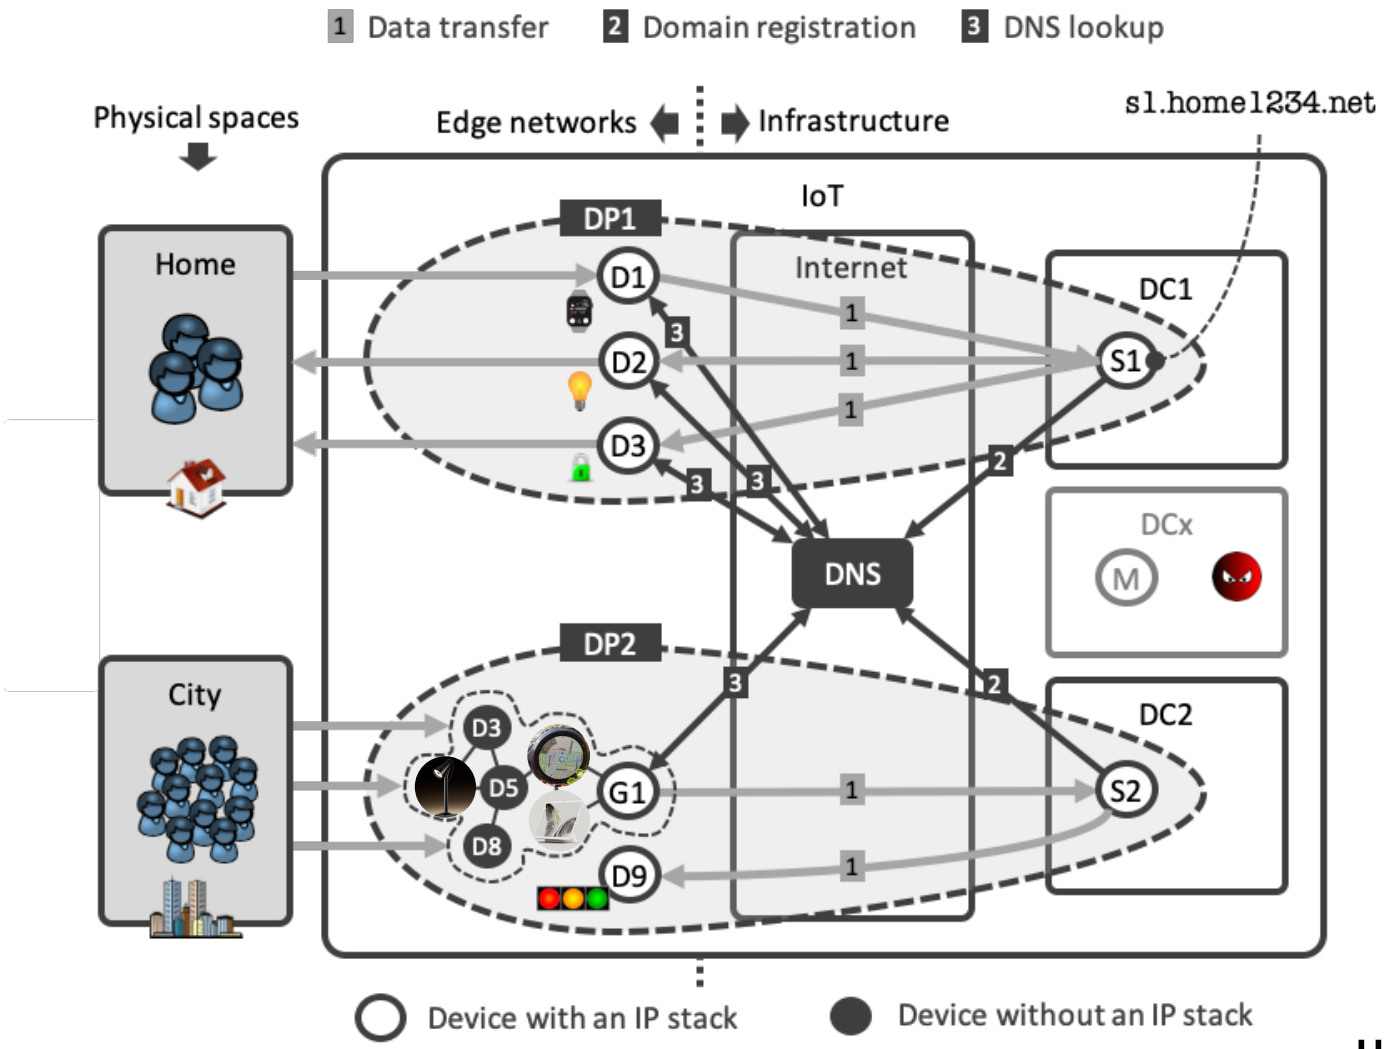

IoT deployments and the Domain Name System (DNS)

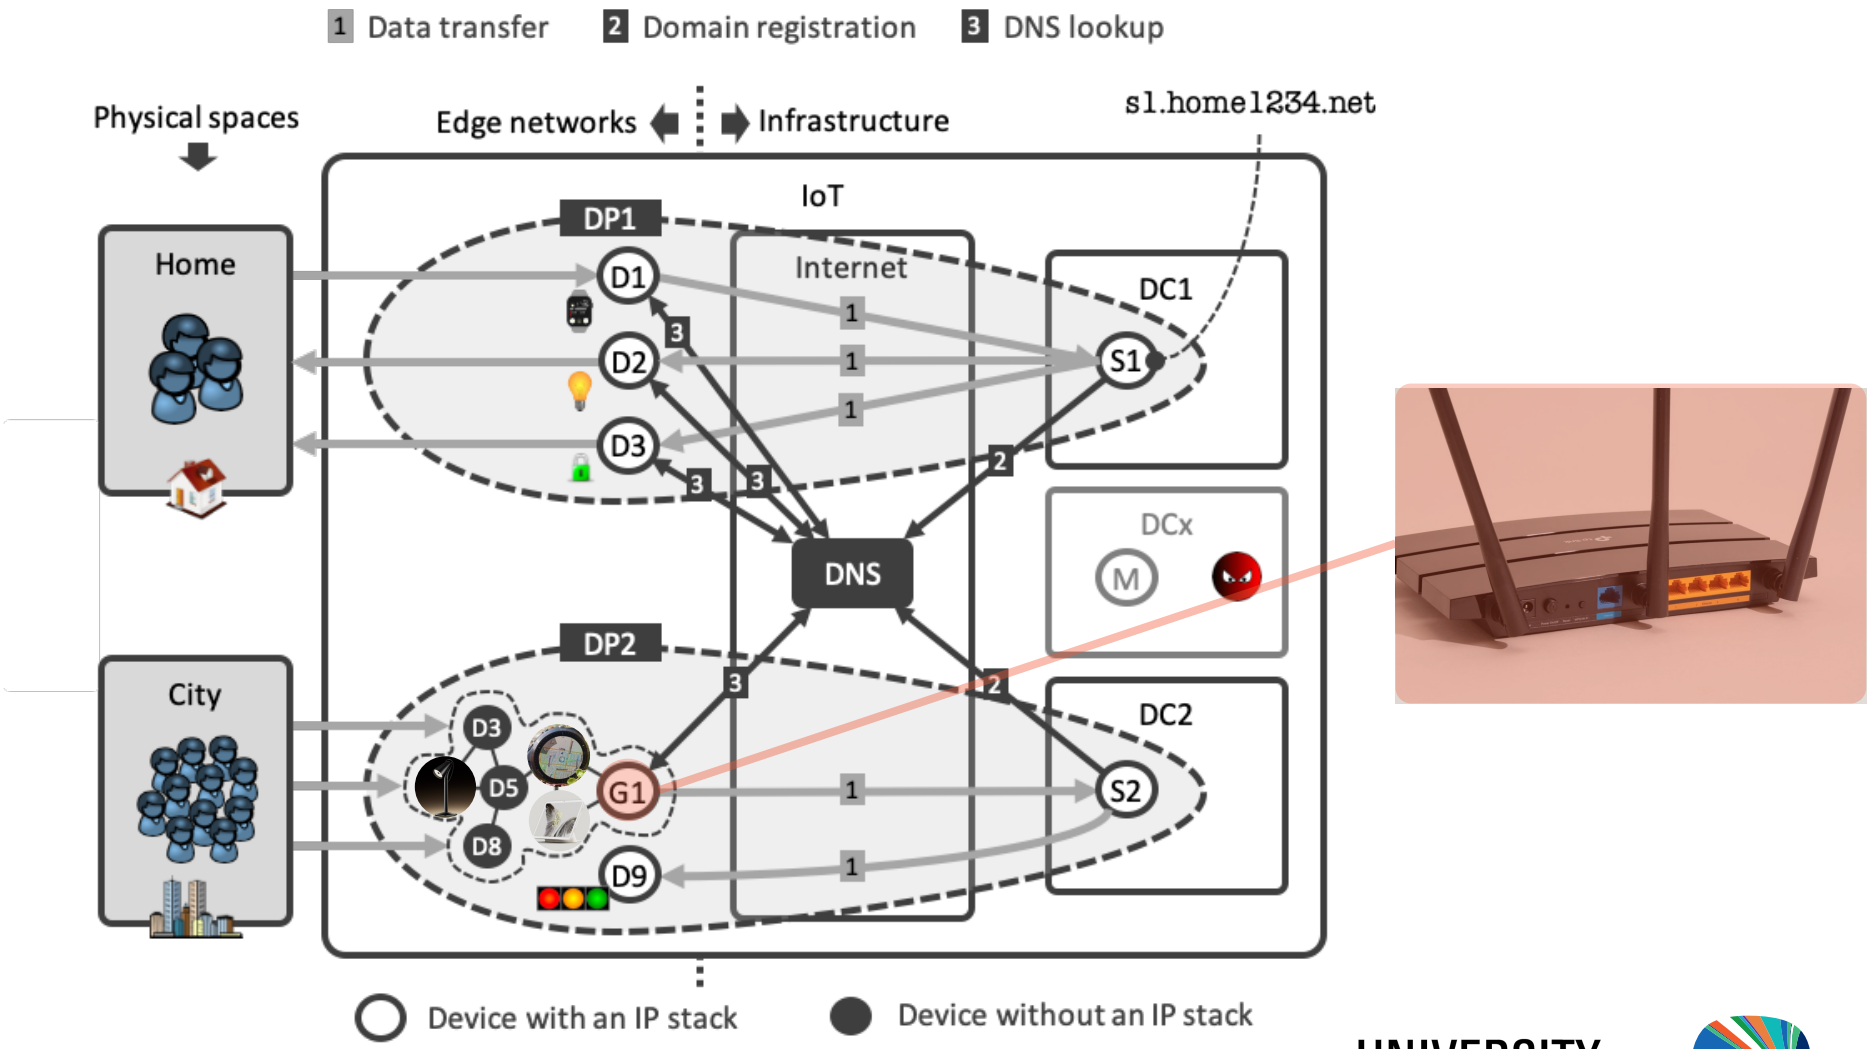

UNIVERSITY OF TWENTE.

SDN LABS

One more thing to check

What is the purpose of DNS caches?

Multiple-choice question:
What's the purpose of DNS caches?

A. Lower DNS response times
B. Increase DNS scalability
C. Enable operators to analyze DNS queries
D. Increase demand for computer memory

Now you know DNS

Let’s look at the current situation

UNIVERSITY OF TWENTE.

SDN LABS

What can we do?

Yes, we have some plans

Overview

Opportunities

O1	Using DoH/DoT to encrypt DNS queries
O2	Using DNSSEC to detect malicious redirects of IoT devices
O3	DNS protocols to double-check the authenticity of IoT services
O4	Protecting IoT devices against domain registration hijacks
O5	Using DNS datasets to increase IoT transparency

Risks

R1	DNS unfriendly programming at IoT scale
R2	Increased size and complexity of IoT botnets targeting the DNS
R3	Increased DDoS amplification through open DNS resolvers

Challenges

C1	Developing a DNS security and transparency library for IoT devices
C2	Training IoT and DNS professionals
C3	Developing a system to share information on IoT botnets
C4	Proactive and flexible mitigation of IoT-powered DDoS traffic
C5	Developing a system to measure how the IoT uses the DNS

UNIVERSITY OF TWENTE.

SIDN LABS

Overview

Opportunities

O1	Using DoH/DoT to encrypt DNS queries
O2	Using DNSSEC to detect malicious redirects of IoT devices
O5	Using DNS datasets to increase IoT transparency

Risks

R1	DNS unfriendly programming at IoT scale
R2	Increased size and complexity of IoT botnets targeting the DNS

Challenges

C1	Developing a DNS security and transparency library for IoT devices
C3	Developing a system to share information on IoT botnets
C4	Proactive and flexible mitigation of IoT-powered DDoS traffic

Overview

Opportunities

Help meet IoT's new safety and transparency requirements

O1 Using DoH/DoT to encrypt DNS queries

O2 Using DNSSEC to detect malicious redirects of IoT devices

O5 Using DNS datasets to increase IoT transparency

Risks

Protect the SSR of the DNS against insecure IoT devices

R1 DNS unfriendly programming at IoT scale

R2 Increased size and complexity of IoT botnets targeting the DNS

Challenges

Technologies and systems that need to be developed

C1 Developing a DNS security and transparency library for IoT devices

C3 Developing a system to share information on IoT botnets

C4 Proactive and flexible mitigation of IoT-powered DDoS traffic

O1 Using DoH/DoT to encrypt DNS queries

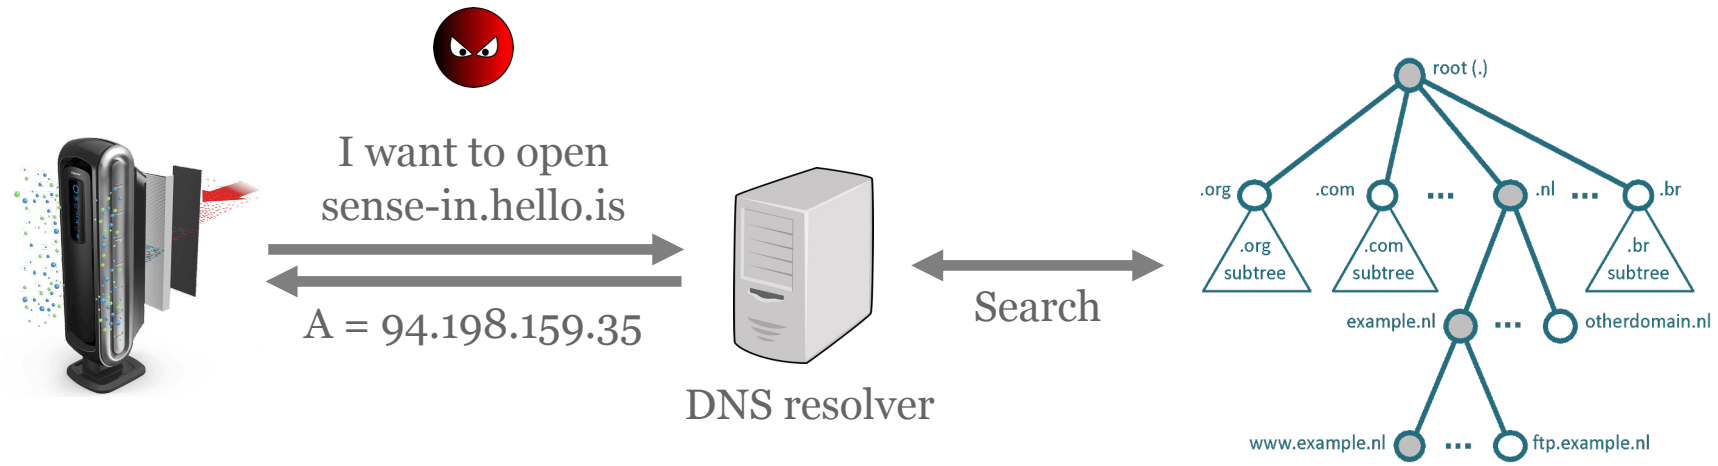

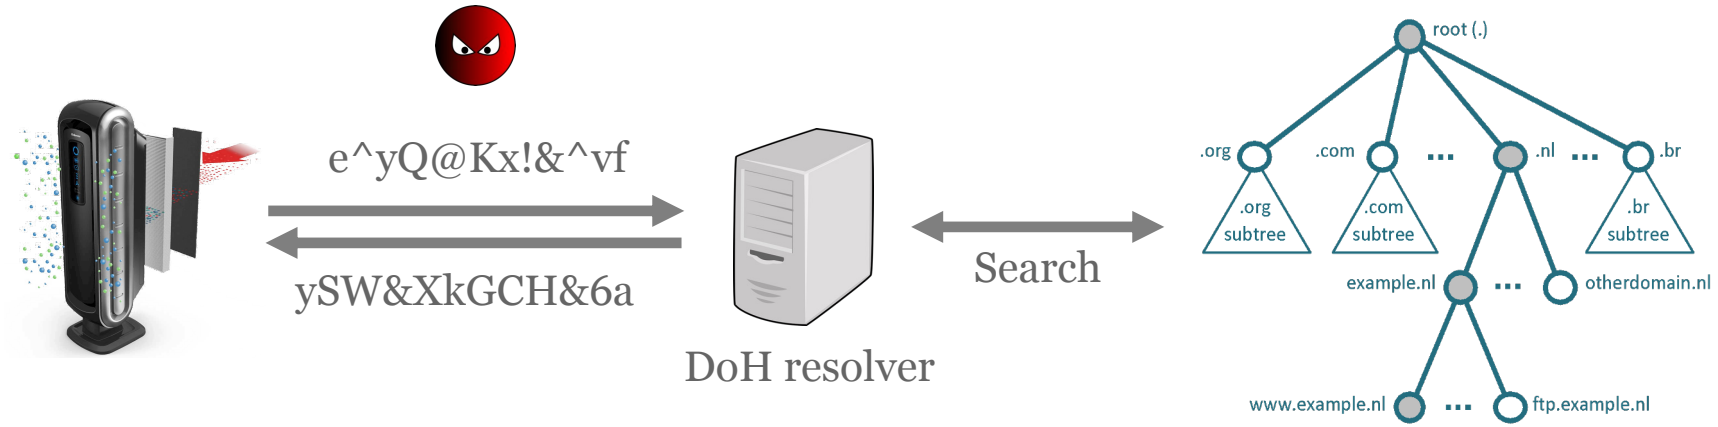

O1 Using DNS-over-HTTPS to encrypt DNS queries

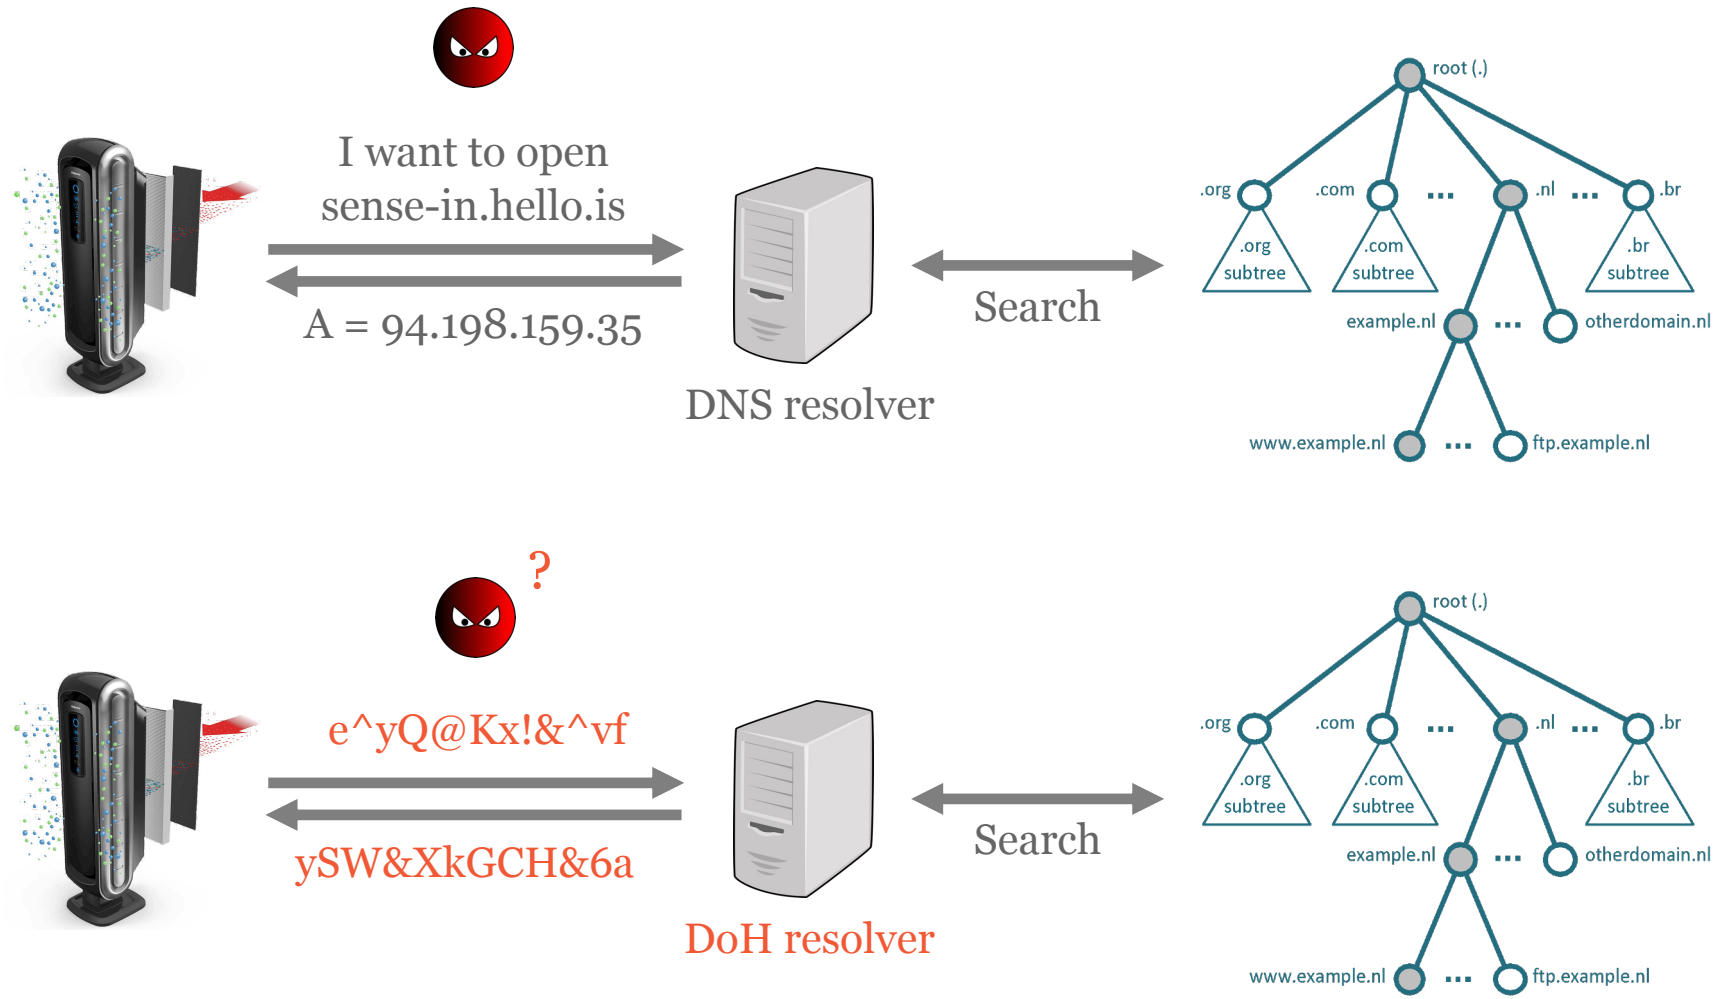

UNIVERSITY OF TWENTE.

SDN LABS

O2 Signing DNS responses with DNSSEC

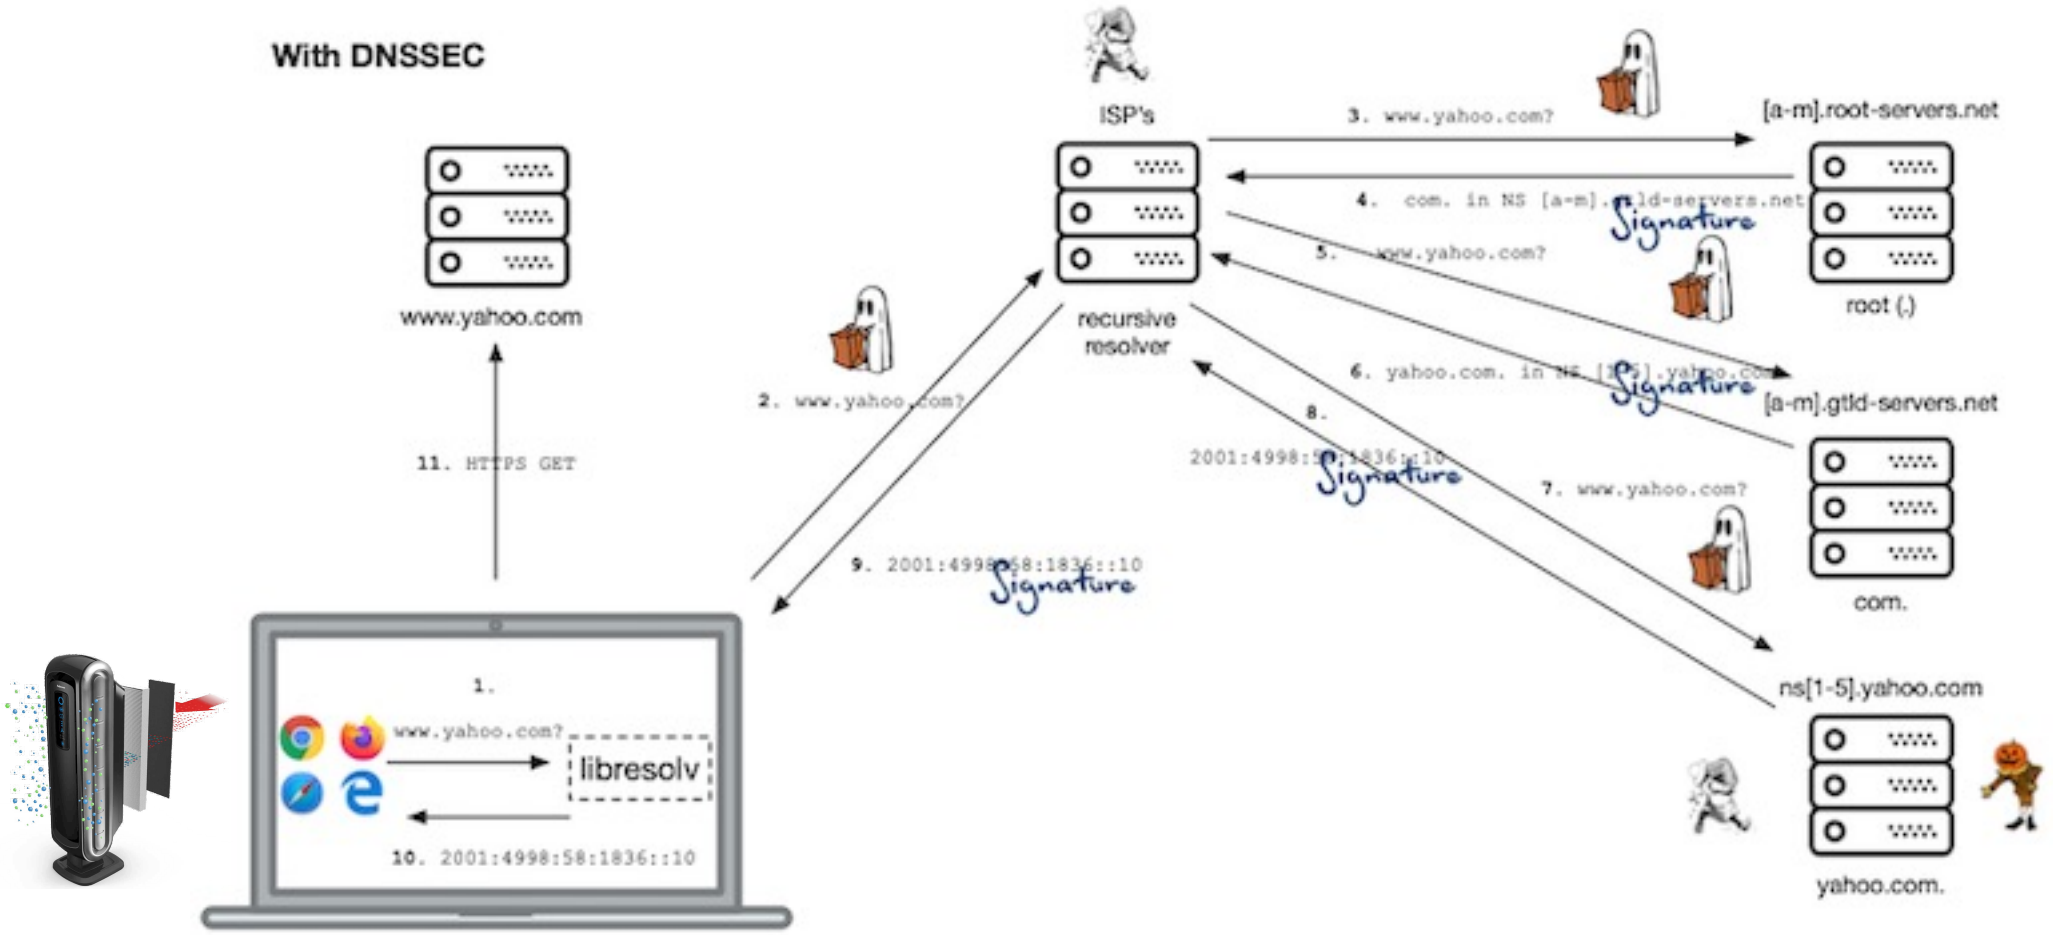

Source: https://www.netmeister.org/blog/doh-dot-dnssec.html

UNIVERSITY OF TWENTE.

SDN LABS

If you’re the IT operators

Would you apply these? Is there still a concern?

UNIVERSITY OF TWENTE.

SDN LABS

O5 Using DNS datasets to increase IoT transparency

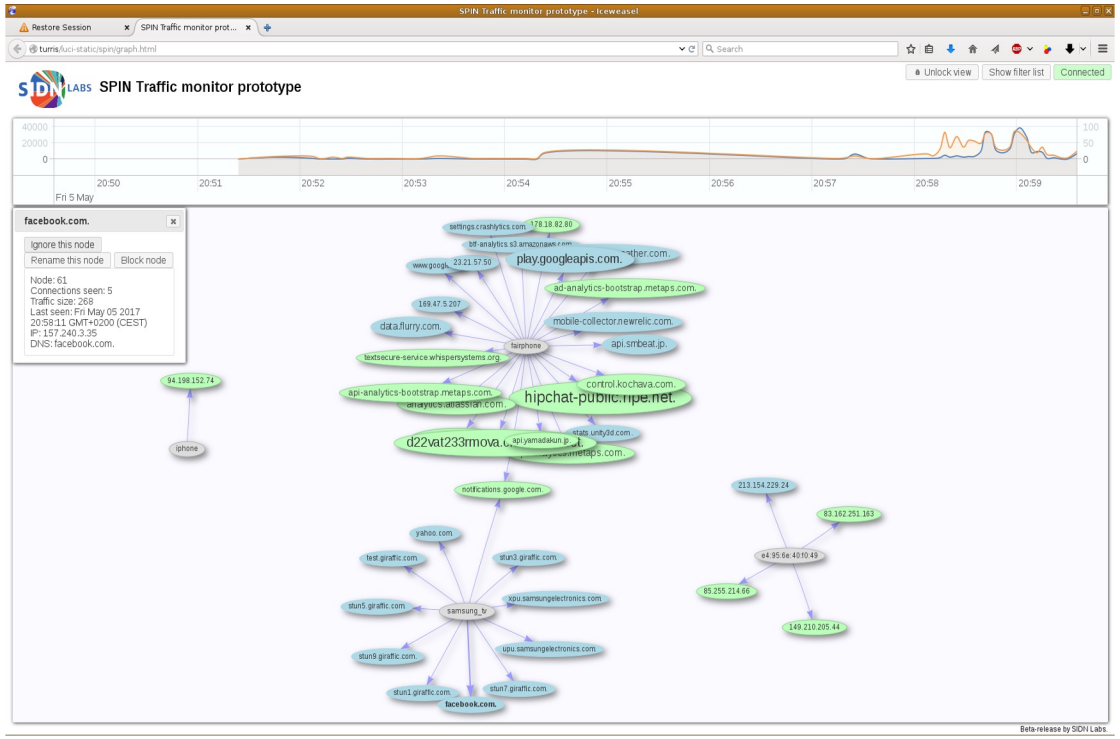

spin.sidnlabs.nl | github.com/sidn/spin

- Measure IoT device's DNS queries
- Requires intuitive visualization for users
- Also, what sensor data are devices sharing?
- Perhaps a topic for future regulation
- Part of larger discussion on data autonomy

UNIVERSITY OF TWENTE.

SIDN LABS

As the end-users

Do you want those graph, data, and regulation?

R1 DNS-unfriendly programming at IoT scale

- TuneIn app example: 700 iPhones generating random queries www.<random-string>.com
- In the stone age (2012), but still: imagine millions of unsupported devices exhibiting that kind of behavior after a software update
- High-level APIs abstract DNS away from developers

- Actually, this does not apply to DNS alone. Unfriendly programming and Software update can cause trouble everywhere like large company

UNIVERSITY OF TWENTE.

SDN LABS

If you’re the manager/engineer

What would you do to prevent this?

R2 DDoS attacks by IoT botnets

- IoT botnets of 400-600K bots (Mirai, Hajime), may increase
- Higher propagation rates (e.g., +50K bots in 24 hours)
- Vulnerabilities difficult to fix, botnet infections unnoticed
- DDoS amplification: 23-25 million open resolvers (now around 3 million)

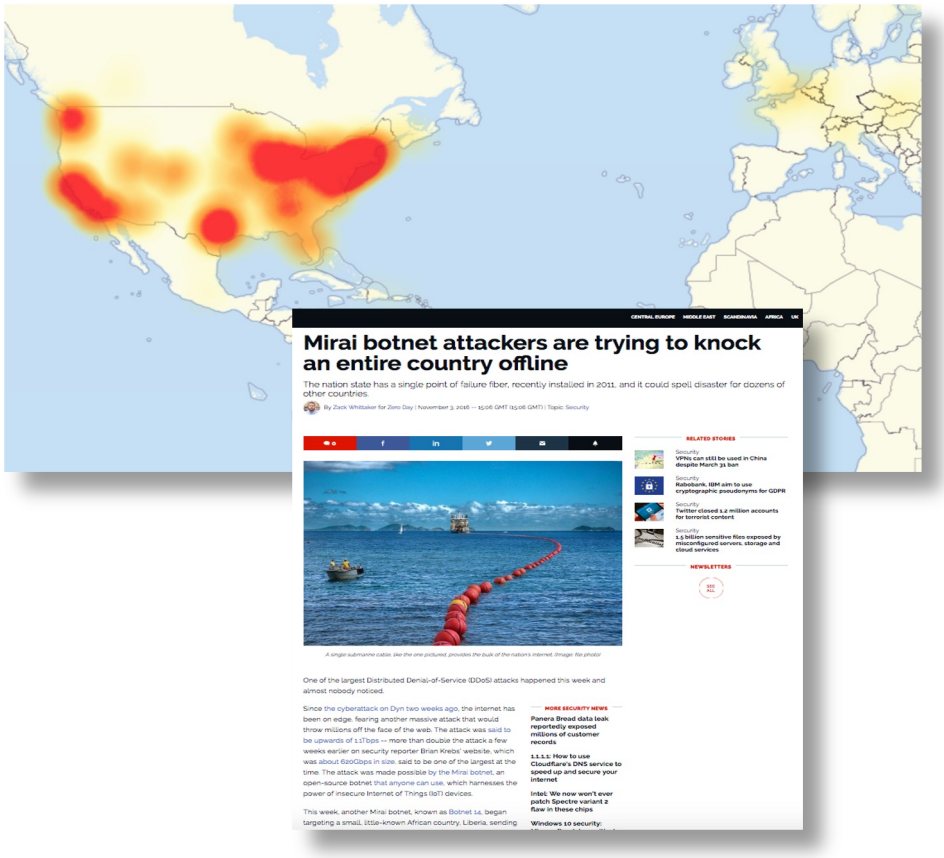

Do you think your device is safe?

What will you do after this lecture?

Open question:
What do you think will make IoT botnets more difficult to eradicate than a traditional ones?

Why collaborative?

- Collaborative incident analysis
- Example: Mirai IoT botnet
- 11 sources, 9 organizations/sites

[Mirai]

Role	Data Source	Collection Site	Collection Period	Data Volume
Growth and size	Network telescope	Merit Network, Inc.	07/18/2016–02/28/2017	370B packets, avg. 269K IPs/min
Device composition	Active scanning	Censys	07/19/2016–02/28/2017	136 IPv4 scans, 5 protocols
Ownership & evolution	Telnet honeypots	AWS EC2	11/02/2016–02/28/2017	141 binaries
	Telnet honeypots	Akamai	11/10/2016–02/13/2017	293 binaries
	Malware repository	VirusTotal	05/24/2016–01/30/2017	594 binaries
	DNS — active	Georgia Tech	08/01/2016–02/28/2017	290M RRs/day
	DNS — passive	Large U.S. ISP	08/01/2016–02/28/2017	209M RRs/day
Attack characterization	C2 milkers	Akamai	09/27/2016–02/28/2017	64.0K attack commands
	DDoS IP addresses	Akamai	09/21/2016	12.3K IP addresses
	DDoS IP addresses	Google Shield	09/25/2016	158.8K IP addresses
	DDoS IP addresses	Dyn	10/21/2016	107.5K IP addresses

Table 1: **Data Sources**—We utilized a multitude of data perspectives to empirically analyze the Mirai botnet.

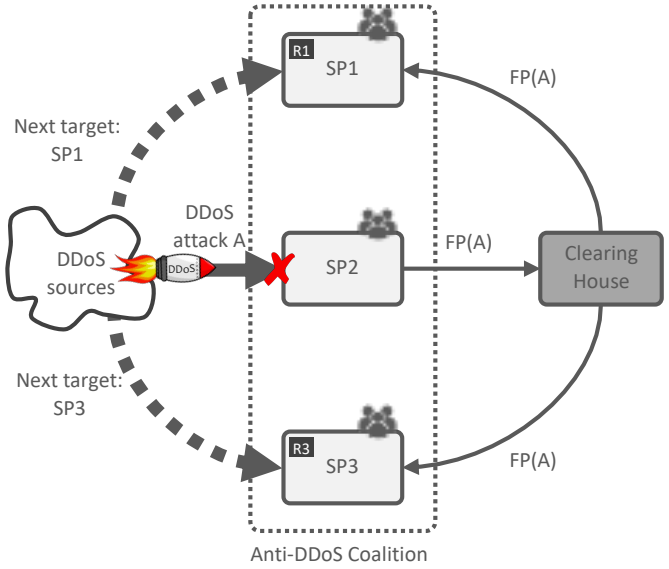

- Collaborative mitigation of (IoT-powered) DDoS attacks
- Fingerprinting of DDoS attacks
- Sharing fingerprints and mitigation rules
- More details: antiddoscoalition.nl

UNIVERSITY OF TWENTE.

SDN LABS

A platform for collaboration

Sounds good, but what are pros and cons?

UNIVERSITY OF TWENTE.

SDN LABS

Challenges for the DNS and IoT industries

- Develop an open-source DNS security and transparency library for IoT devices
  - Such as DNSSEC validation, DoH/DoT support
  - User control over DNS security settings and services used

- Develop a system to proactively detect IoT botnets
  - Share DDoS “fingerprints”, countermeasures, and other botnet characteristics across operators
  - **Collaborative** DDoS detection and learning

- **Collaboratively** handle IoT-powered DDoS attacks
  - DDoS mitigation broker to flexibly share mitigation capacity
  - Security systems in edge networks, such as home routers

Key takeaways

- IoT enables smarter, safer, more sustainable society, but extraordinary safety and privacy risks
- The DNS is one of the core components of the Internet infrastructure for traditional applications and will also play a key role for the IoT
- Opportunities to help fulfilling the IoT's new safety and transparency requirements using the DNS' security functions, datasets, and ubiquitous nature
- Poorly developed and maintained IoT devices are a risk in terms of security and DNS usage
- Many challenges for the interaction between the IoT and the DNS, but starting points exist

UNIVERSITY OF TWENTE.

SIDN LABS

You need to know your enemies

UNIVERSITY OF TWENTE.

SDN LABS

Now you're ready

What would you say when people ask?

UNIVERSITY OF TWENTE.

SDN LABS

“Illuminating Large-Scale IPv6 Scanning in the Internet”

22nd ACM Internet Measurement Conference (IMC ’22), New York, NY, USA, 410–418, 2022,

UNIVERSITY OF TWENTE.

SIDN LABS

Learning Goals

- To understand challenges of IPv6 scanning and scan detection

- To become familiar with common scanning practices in IPv6 in the wild

Quiz Question

How long would it take to scan the **IPv4** address space on a typical desktop computer with a gigabit Ethernet connection, approximately?

A. A week
B. A day
C. An hour
D. A minute

Have you already experimented with Internet-wide scans?

How long would it take to scan IPv6?

Discussion Question #1

- How would you scan IPv6?
- How would your scanning infrastructure look like?

IoT Botnets

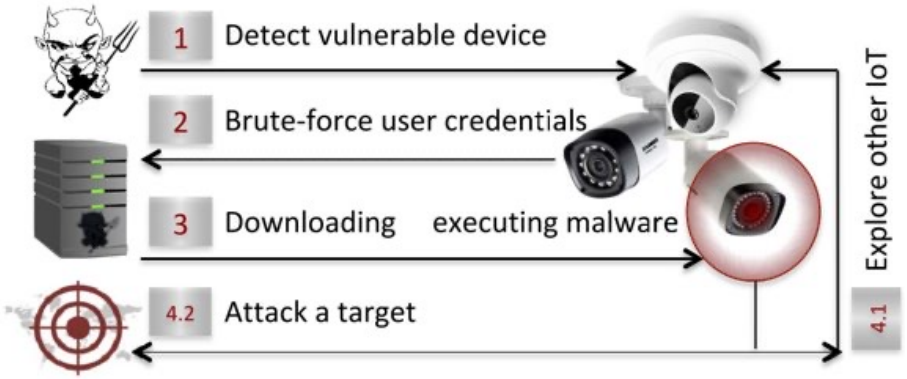

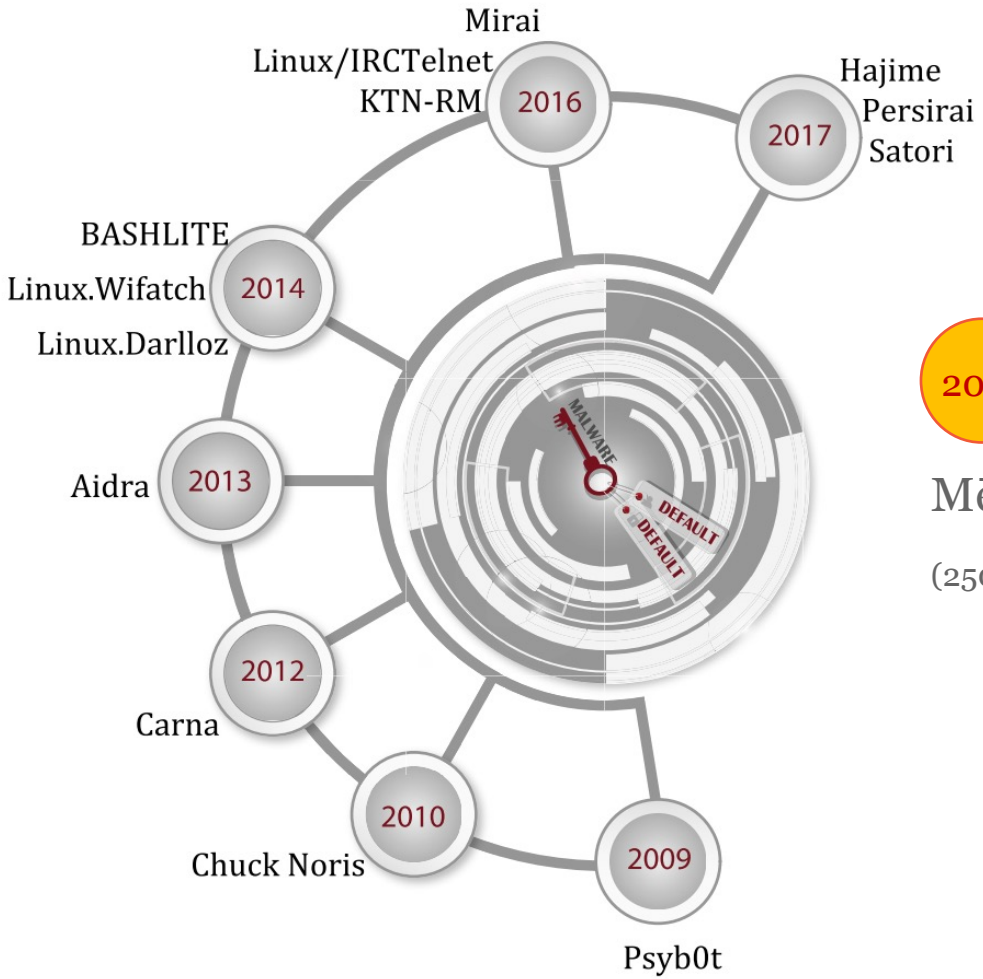

Figures from: Neshenko et al., Demystifying IoT Security: An Exhaustive Survey on IoT Vulnerabilities and a First Empirical Look on Internet-Scale IoT Exploitations

UNIVERSITY OF TWENTE.

SDN LABS

Full IPv6 Scanning

- Using the current rates of IPv4 scans, it would take

$9*10^{24}$ years[1]

to run a full IPv6 scan[2].

- Not even scalable if we use all IoT devices[2] in the world to conduct the scan!

1) $2^{128}/(2^{32}*24*365)$

2) This includes reserved ranges as well, which are not typical scan targets.

3) Estimated to be 20B~30B

Allocated IPv6 Scanning

How long would it take to scan the already allocated IPv6 address space?

Currently* 2344177 /32s are allocated.

(2^96)*2344177 ≈ 1.86*10^35 individual IPs

Still would take 5*10^21 years to scan!

Next Step to reduce our search space?

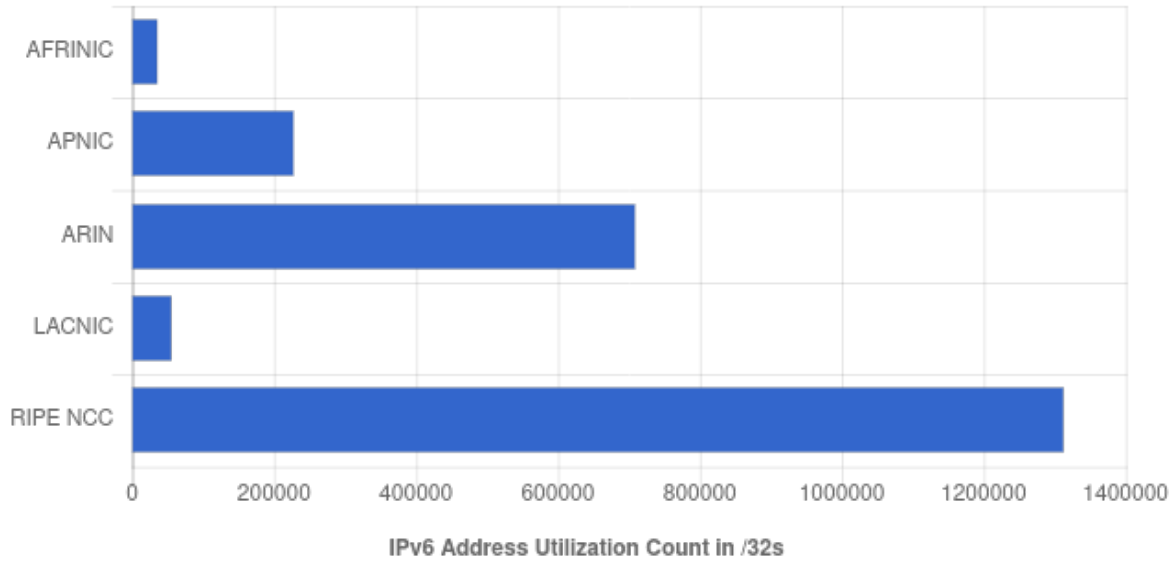

Source: https://www.iana.org/numbers/allocations/

* On 2023-May-02

UNIVERSITY OF TWENTE.

SDN LABS

Target Addresses

- Authors investigate forward DNS entries: 75% of the /64s only target addresses in DNS.
- How would you create an IPv6 hitlist?
- The paper proposes using DNS records and then scanning other nearby addresses (this doesn't hold for all scanners, though).

IPv6 hitlists

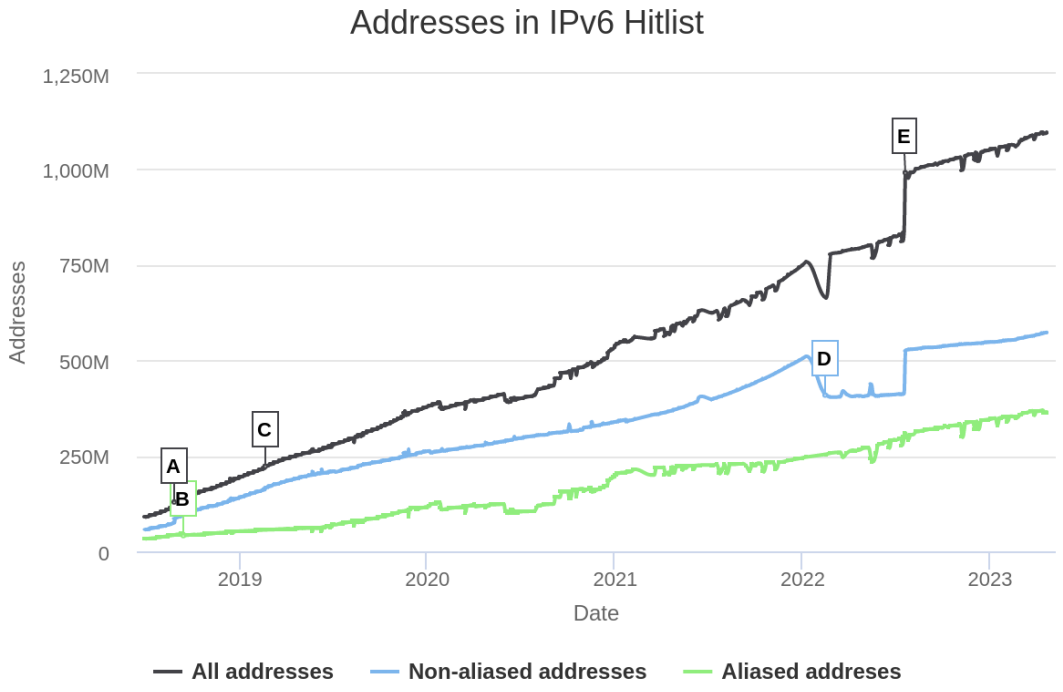

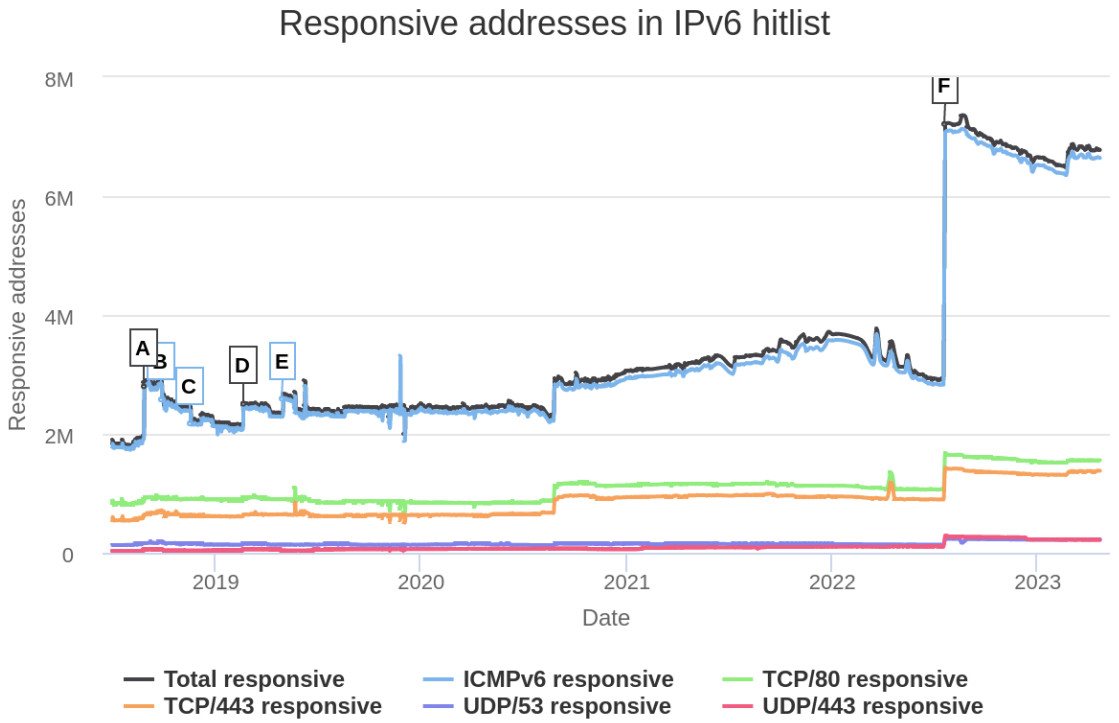

https://ipv6hitlist.github.io/

UNIVERSITY OF TWENTE.

SDN LABS

Additional Reading

- O. Gasser et al., "Scanning the IPv6 Internet: Towards a Comprehensive Hitlist", TMA 2016.
- O. Gasser et al., "Clusters in the Expanse: Understanding and Unbiasing IPv6 Hitlists", IMC 2018.
- J. Zirngibl et al., "Rusty Clusters? Dusting an IPv6 Research Foundation", IMC 2022.

Scan Detection

- How were IPv6 scanners detected in the paper?
  - Authors leverage a CDN network (can you guess?) which is not always feasible.
  - They only consider sources that contact 100 destination IPs with a timeout of 3600 seconds.
  - Probes that hit the same IP more than 5 times are removed.
  - Ports 80 and 443 are not considered (lots of legitimate use), what if IPv6 scans become more port specific?

Discussion Question #2

- How would you detect IPv6 scanners?
  - Detection vantage points
  - Aggregation level (too coarse: conflating individual scan actors, too specific: can lead to missing scanning activities in part or entirely)
  - Other design choices?
- What would be a sound IDS policy to block IPv6 scanners? Can we have an adaptive aggregation?

Scan Sources

- The top-10 source ASes account for more than 99% of scan packets.
- Scans in IPv6 are mostly limited to datacenters and cloud providers. No exclusively residential ISPs in the top 20.
- What else do you find interesting from these two tables?

aggregation	scans	packets	sources	ASes
/128	65,485	2.04B	3,542	55
/64	5,199	2.14B	1,326	62
/48	5,019	2.15B	1,372	76

Table 1: Detected scans over the course of our measurement window (Jan 2021 until Mar 2022). Depending on the aggregation of source IP addresses, the number of scans and scan sources changes dramatically.

			scan sources		
rank	AS type	packets	/48s	/64s	/128s
#1	Datacenter (CN)	839M (39.2%)	1	1	1
#2	Datacenter (CN)	744M (34.8%)	1	1	5
#3	Cybersecurity (US)	275M (12.9%)	1	1	12
#4	Cloud (US/global)	78M (3.7%)	2	2	512
#5	Cloud (DE)	48M (2.3%)	3	59	59
#6	Cloud (US/global)	45M (2.1%)	10	15	205
#7	Cloud (US/global)	39M (1.8%)	9	9	123
#8	Cloud (CN)	30M (1.4%)	5	5	53
#9	Transit (global)	11M (0.5%)	1	2	956
#10	Cloud (CN)	10M (0.5%)	1	1	7
#11	Cloud (US/global)	4.7M (0.2%)	1	1	353
#12	Datacenter (CN)	3.1M (0.1%)	9	12	19
#13	ISP (VN)	2.5M (0.1%)	1	1	1
#14	Datacenter (CN)	1.6M (≤ 0.1%)	1	1	2
#15	Research (DE)	1.1M (≤ 0.1%)	1	1	1
#16	ISP (RU)	0.9M (≤ 0.1%)	1	1	2
#17	University (DE)	0.8M (≤ 0.1%)	1	1	2
#18	Cloud/Transit (DE)	0.6M (≤ 0.1%)	1,092	1,057	1,057
#19	ISP (RU)	0.6M (≤ 0.1%)	1	1	1
#20	University (DE)	0.5M (≤ 0.1%)	1	1	1

UNIVERSITY OF TWENTE.

SIDN LABS

Target Ports

- IPv6 scans currently scan a range of ports similar to penetration testing (IPv4 scans typically target a single port).
  - AS #1 targets some 444 different ports in the first half of 2021, and then only ports 22, 3389, 8080, and 8443 starting in May 2021.
  - AS #3: almost the entire port space, 45k ports.
  - AS #18: only scans port 22.
- Port selection characteristics can be used to attribute scans to entities.
- Which ports would you scan?

Key Takeaways

- IPv6 not only makes scanning more complicated, but also challenges scan detection efforts.
- IPv6 scanners target a broad range of ports, in contrast to IPv4 scans.
- IPv6 scanning is presumably not yet originating from IoT botnets.

Today's learning objective revisited

To what extent to you think you'll be able to discuss the correlation between IoT security and Internet core protocols?

UNIVERSITY OF TWENTE.

SDN LABS

Volg ons

.nl SIDN.nl

@SIDN

SIDN

Q&A

Next lecture: **Wed May 10, 10:45-12:30**

Cristian Hesselman
Director of SIDN Labs
+31 6 25 07 87 33
c.e.w.hesselman@utwente.nl
@hesselma

Elmer Lastdrager
Research Engineer
+31 6 12 47 84 88
elmer.lastdrager@sidn.nl
@ElmerLastdrager

UNIVERSITY OF TWENTE.

SIDN LABS