Lecture #5: IoT Honeypots

Cristian Hesselman, Elmer Lastdrager, Ramin Yazdani, and Etienne Khan

University of Twente | May 20, 2020

UNIVERSITY OF TWENTE.

SIDN LABS

Lab assignment

- MUD descriptions: you'll need to generate them yourselves, tools are available
- IoT devices: you'll need to work with the actual hardware, no emulations (unless as an extra)
- Use IoT devices without a browser-like interface, such as light bulbs, audio speakers, doorbells
- Do not use multi-purpose devices like tablets, phones, laptops
- At least 2 IoT devices per group of 3 and at least 3 devices per group of 4
- Etienne Khan available for assistance

UNIVERSITY OF TWENTE.

SDN LABS

Paper summaries

- You must have handed in your two summaries BEFORE this lecture
- You can use the summaries during the oral exam (“open book”)
- You **cannot** complete SSI without submitting 12 paper summaries!

Interactive Lecture

- Goal: enable you to learn from each other and further increase your understanding of the papers (contributes to preparing yourself for the oral exam)

- Format:
  1. We'll ask someone to provide their verbal summary of the paper
  2. 5-slide(-ish) summary by teachers (put any questions in the chat)
  3. Questions: discussion starters and fact questions
  4. Discussion (use your mic)
  5. We may ask someone specific to start the discussion

- Experimental format resulting from Corona pandemic, please provide feedback!

UNIVERSITY OF TWENTE.

SDN LABS

Today's papers

Are about measuring IoT botnets

- **[IoTPOT]** Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, Christian Rossow. "IoTPOT: Analysing the Rise of IoT Compromises". 9th USENIX Workshop on Offensive Technologies (co-located with USENIX Sec '15), WOOT '15, Washington, DC, https://christian-rossow.de/publications/iotpot-woot2015.pdf

- [**Honware**] Vetterl, Alexander, and Richard Clayton. "Honware: A virtual honeypot framework for capturing CPE and IoT zero days." Symposium on Electronic Crime Research (eCrime). IEEE. 2019. https://www.cl.cam.ac.uk/~amv42/papers/vetterl-clayton-honware-virtual-honeypot-framework-ecrime-19.pdf

“IoTPOT: Analysing the Rise of IoT Compromises”, 9th USENIX Workshop on Offensive Technologies (WOOT), 2015

UNIVERSITY OF TWENTE.

SDN LABS

Darknet monitoring

270.000 IP's

Connect back 23/80 TCP

& collect banners.

Table 1 - Scanning hosts and device models

Device Type	Host Count	Device Model Count
DVR	1,509	19
IP Camera	523	16
Wireless Router	118	45
Customer Premises Equipment	65	1
Industrial Video Server	22	1
TV Receiver	19	2
Heat Pump	10	1
EMU System	9	1
Digital Video Scalar	5	2
Router	4	3

Darknet monitoring (2)

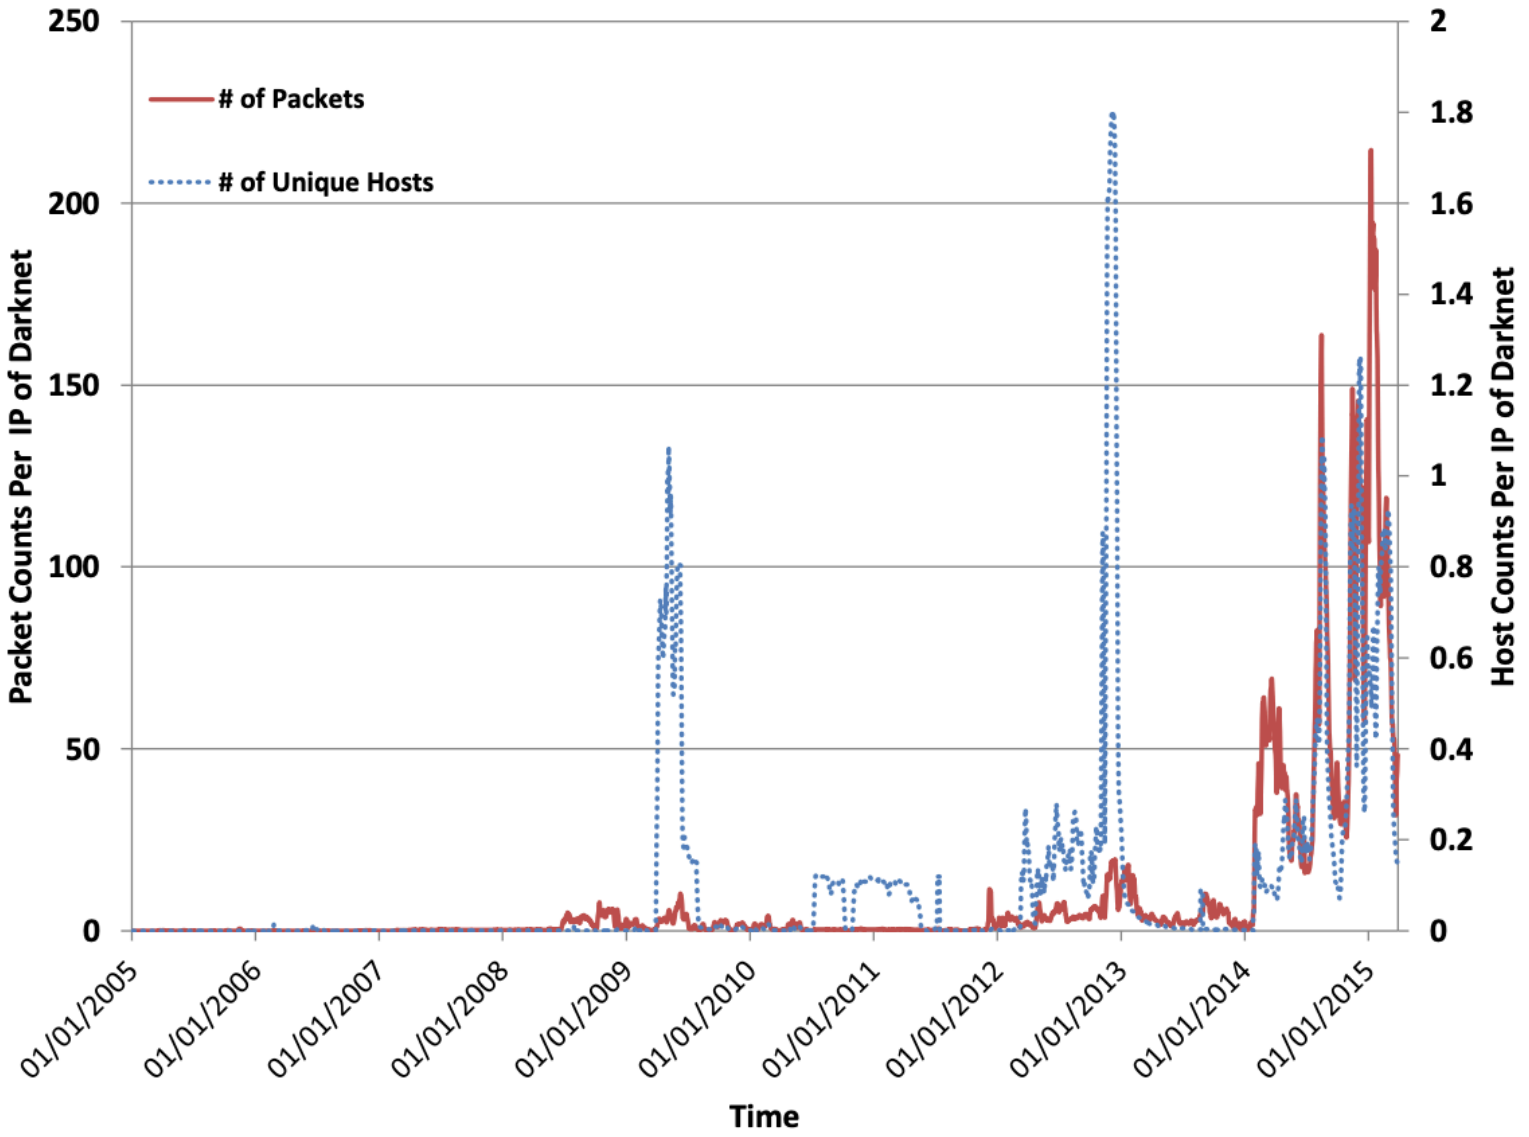

UNIVERSITY OF TWENTE.

SDN LABS

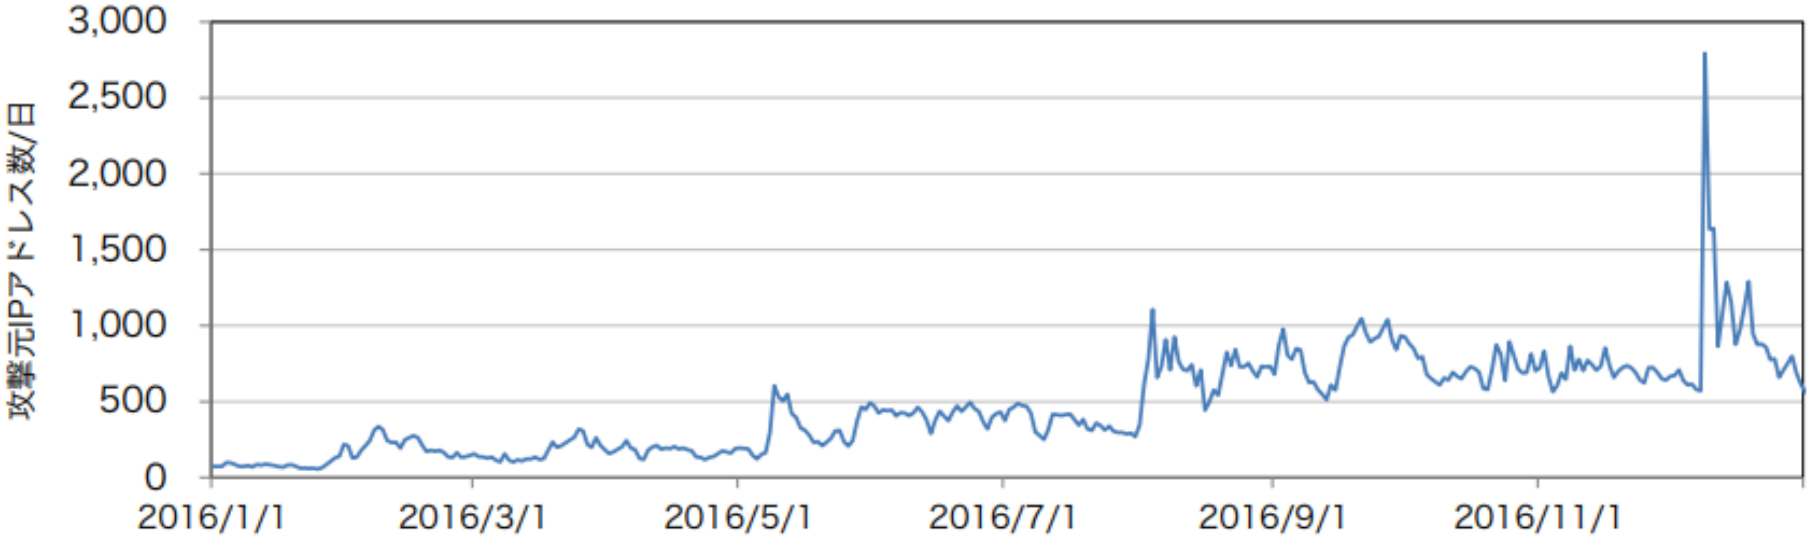

図5: 23/TCP に対する日本国内からの攻撃

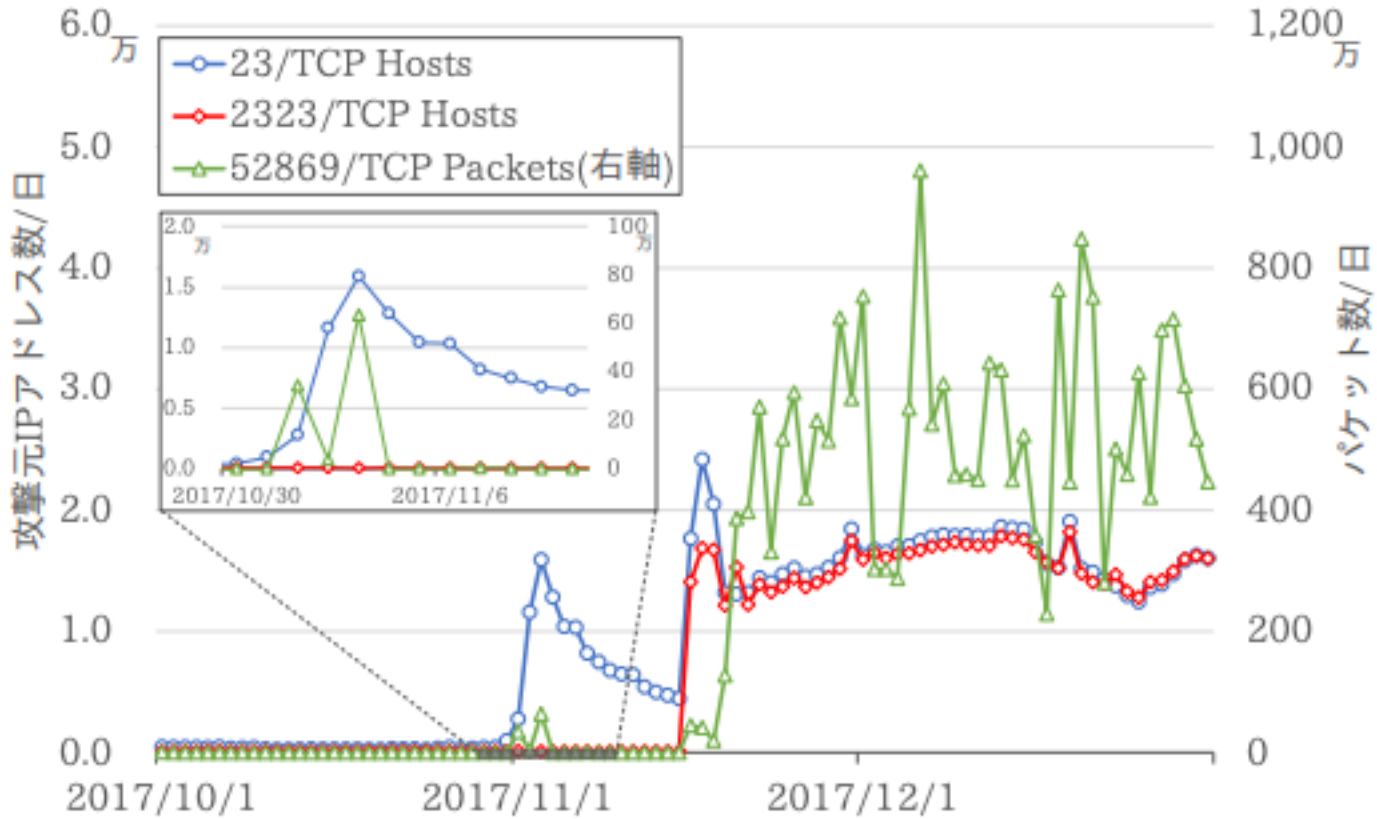

Quiz

Why is a **darknet** useful for IoT malware research?

A: Malware runs better, because it's from the dark side

B: No legitimate traffic

C: No legal problems because a darknet is not managed by any company

D: It has residual trust from previous use

IoT POT

Running on 165 IP addresses

5 weeks running time

Telnet attack stages:

(1) Intrusion; (2) Infection; (3) Monetization. *Remember Mirai?*

Credentials in Fixed/Random order (1)

6 patterns of commands (2) distinguished

‘Coordinated intrusion’

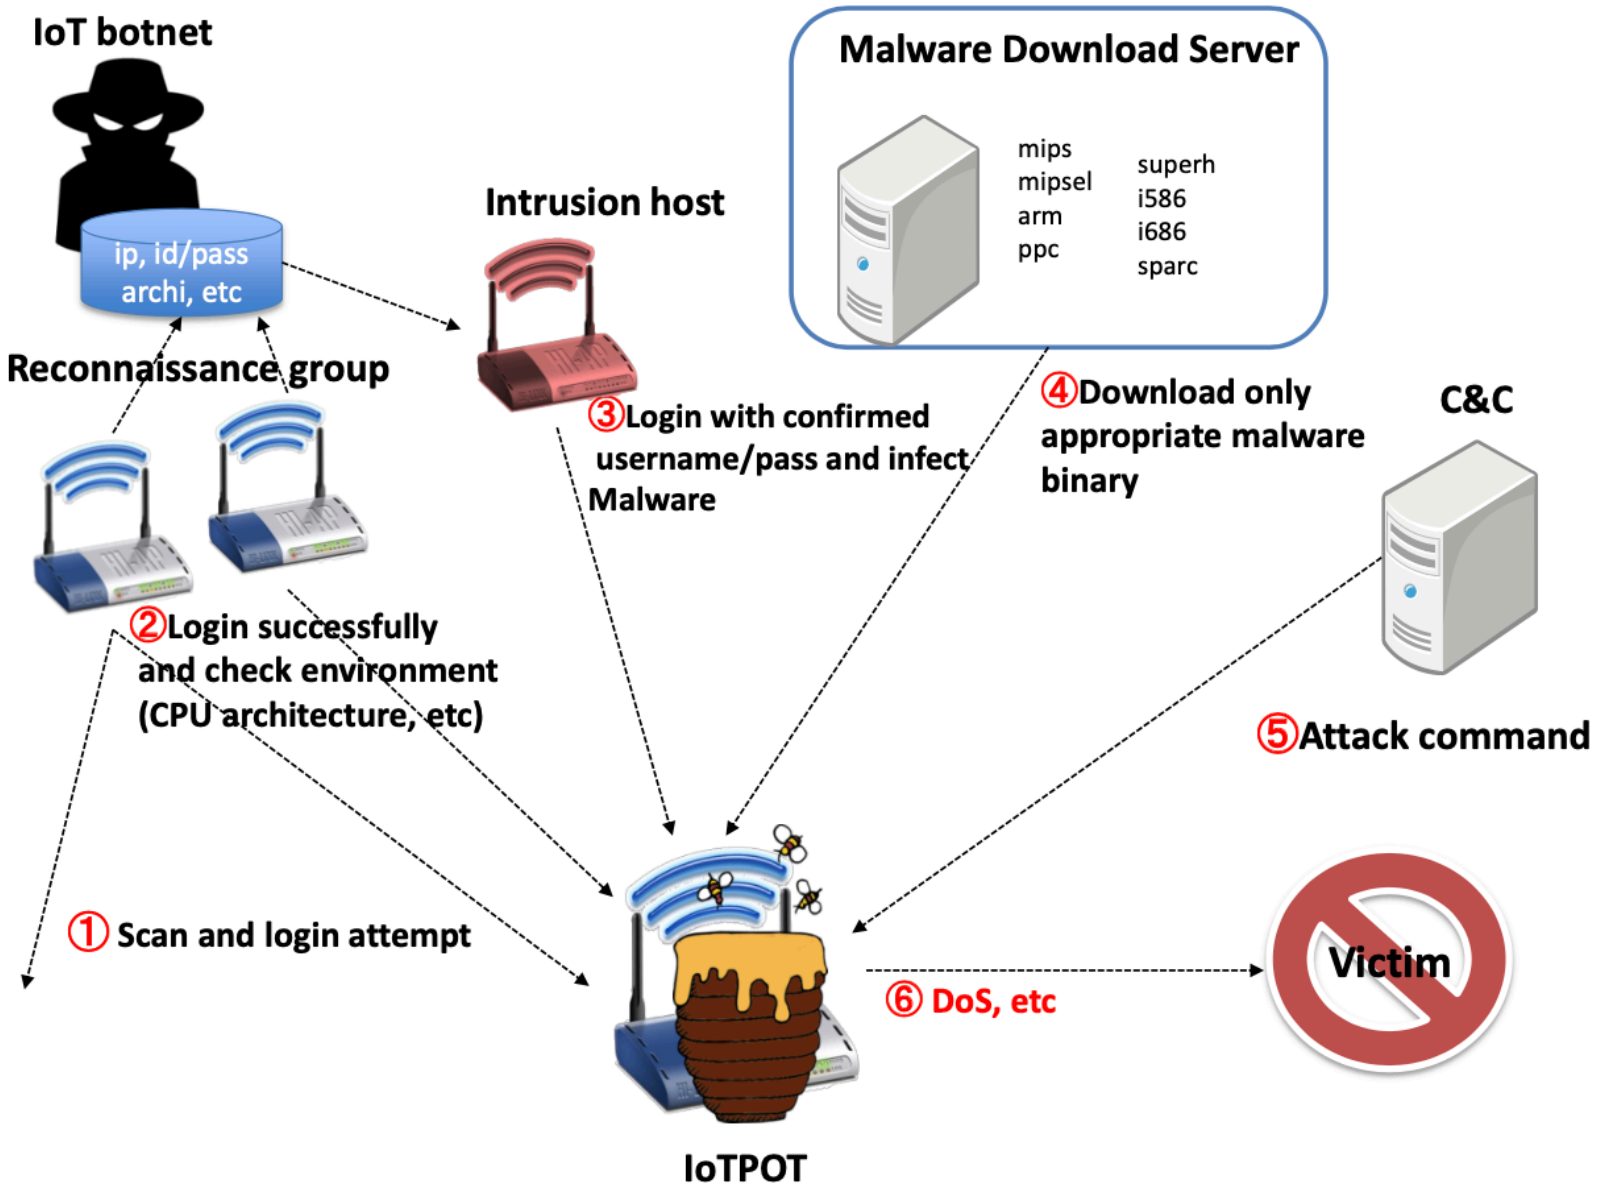

UNIVERSITY OF TWENTE.

SDN LABS

IoTPOT & IoTBOX

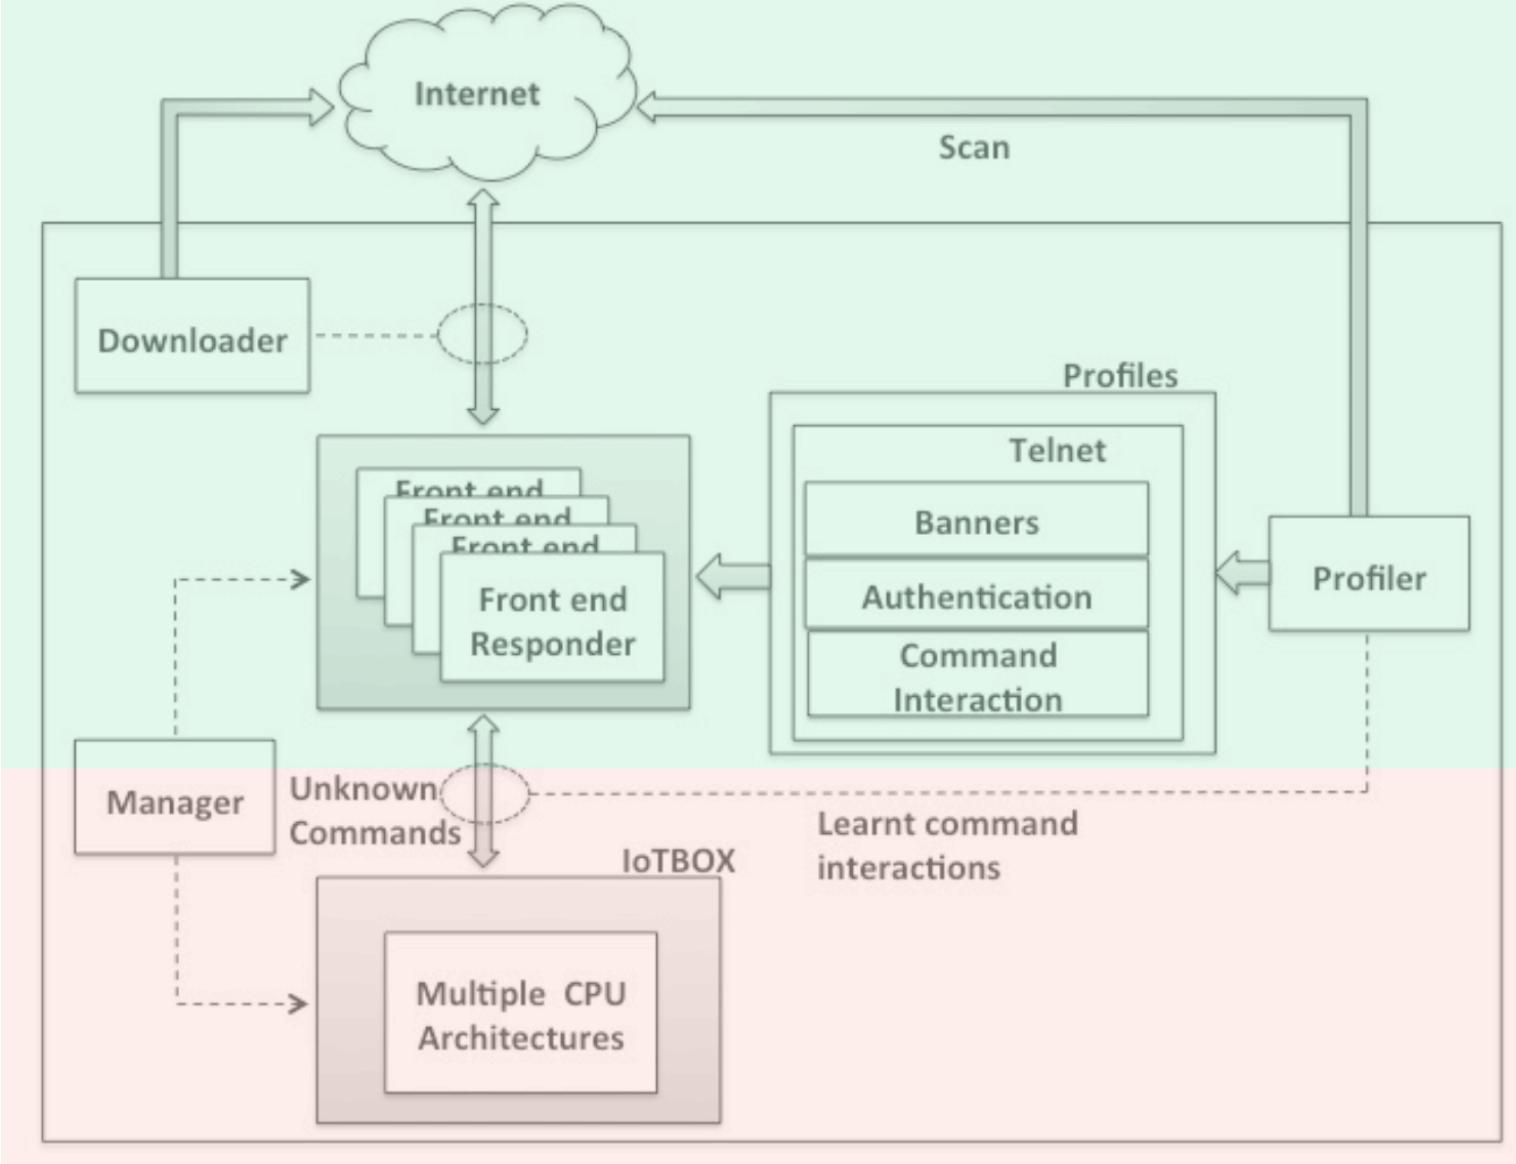

UNIVERSITY OF TWENTE.

SDN LABS

Quiz

What would an operator of an IoTPOT honeypot need to do to support Hajime?

A: Add support for MIPS CPU architecture

B: Track DHT (P2P) communications

C: Expose many vulnerabilities

D: Run the honeypot in different subnets

IoTBOX

Sandbox with 8 CPU architectures

Limit outgoing to DNS/HTTP 5ppm

Telnet to Dummy server

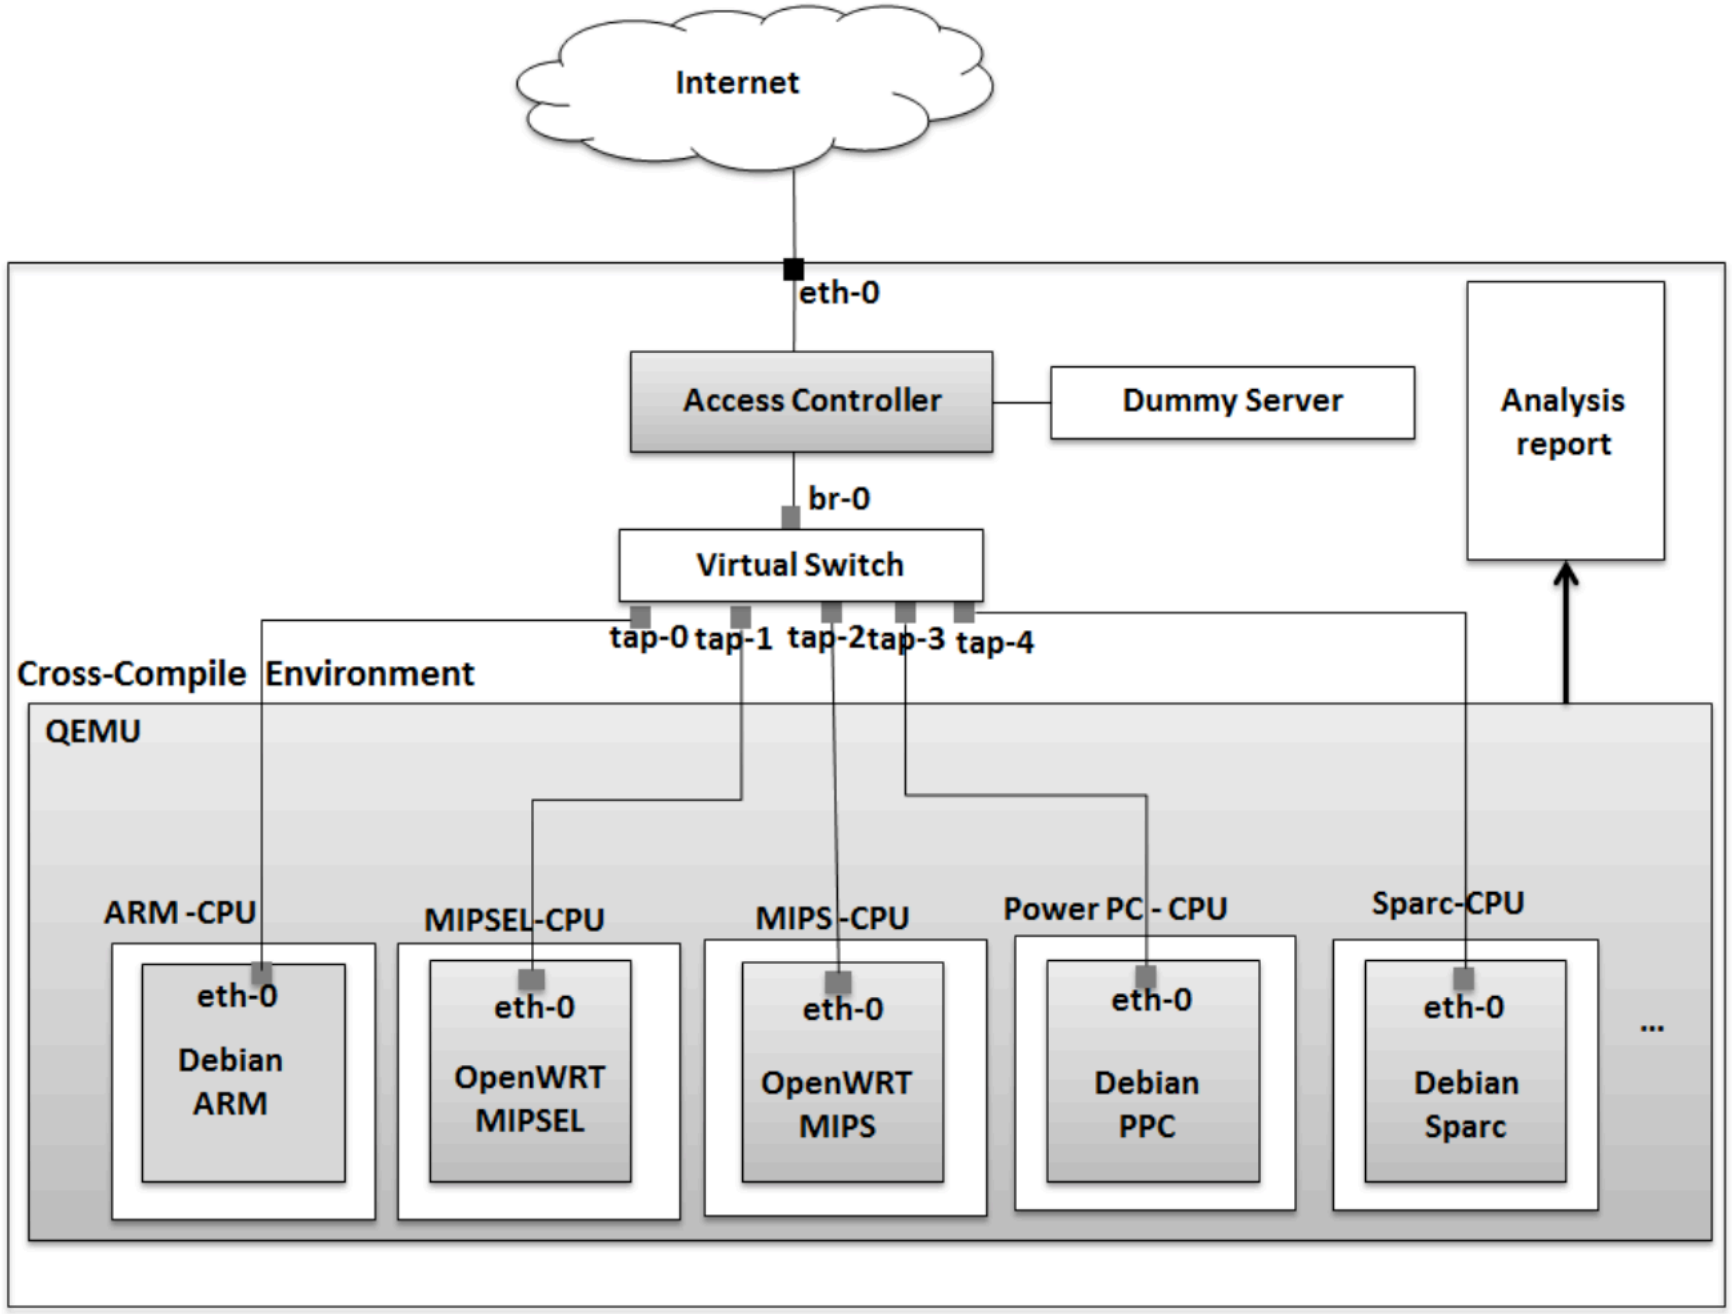

Results

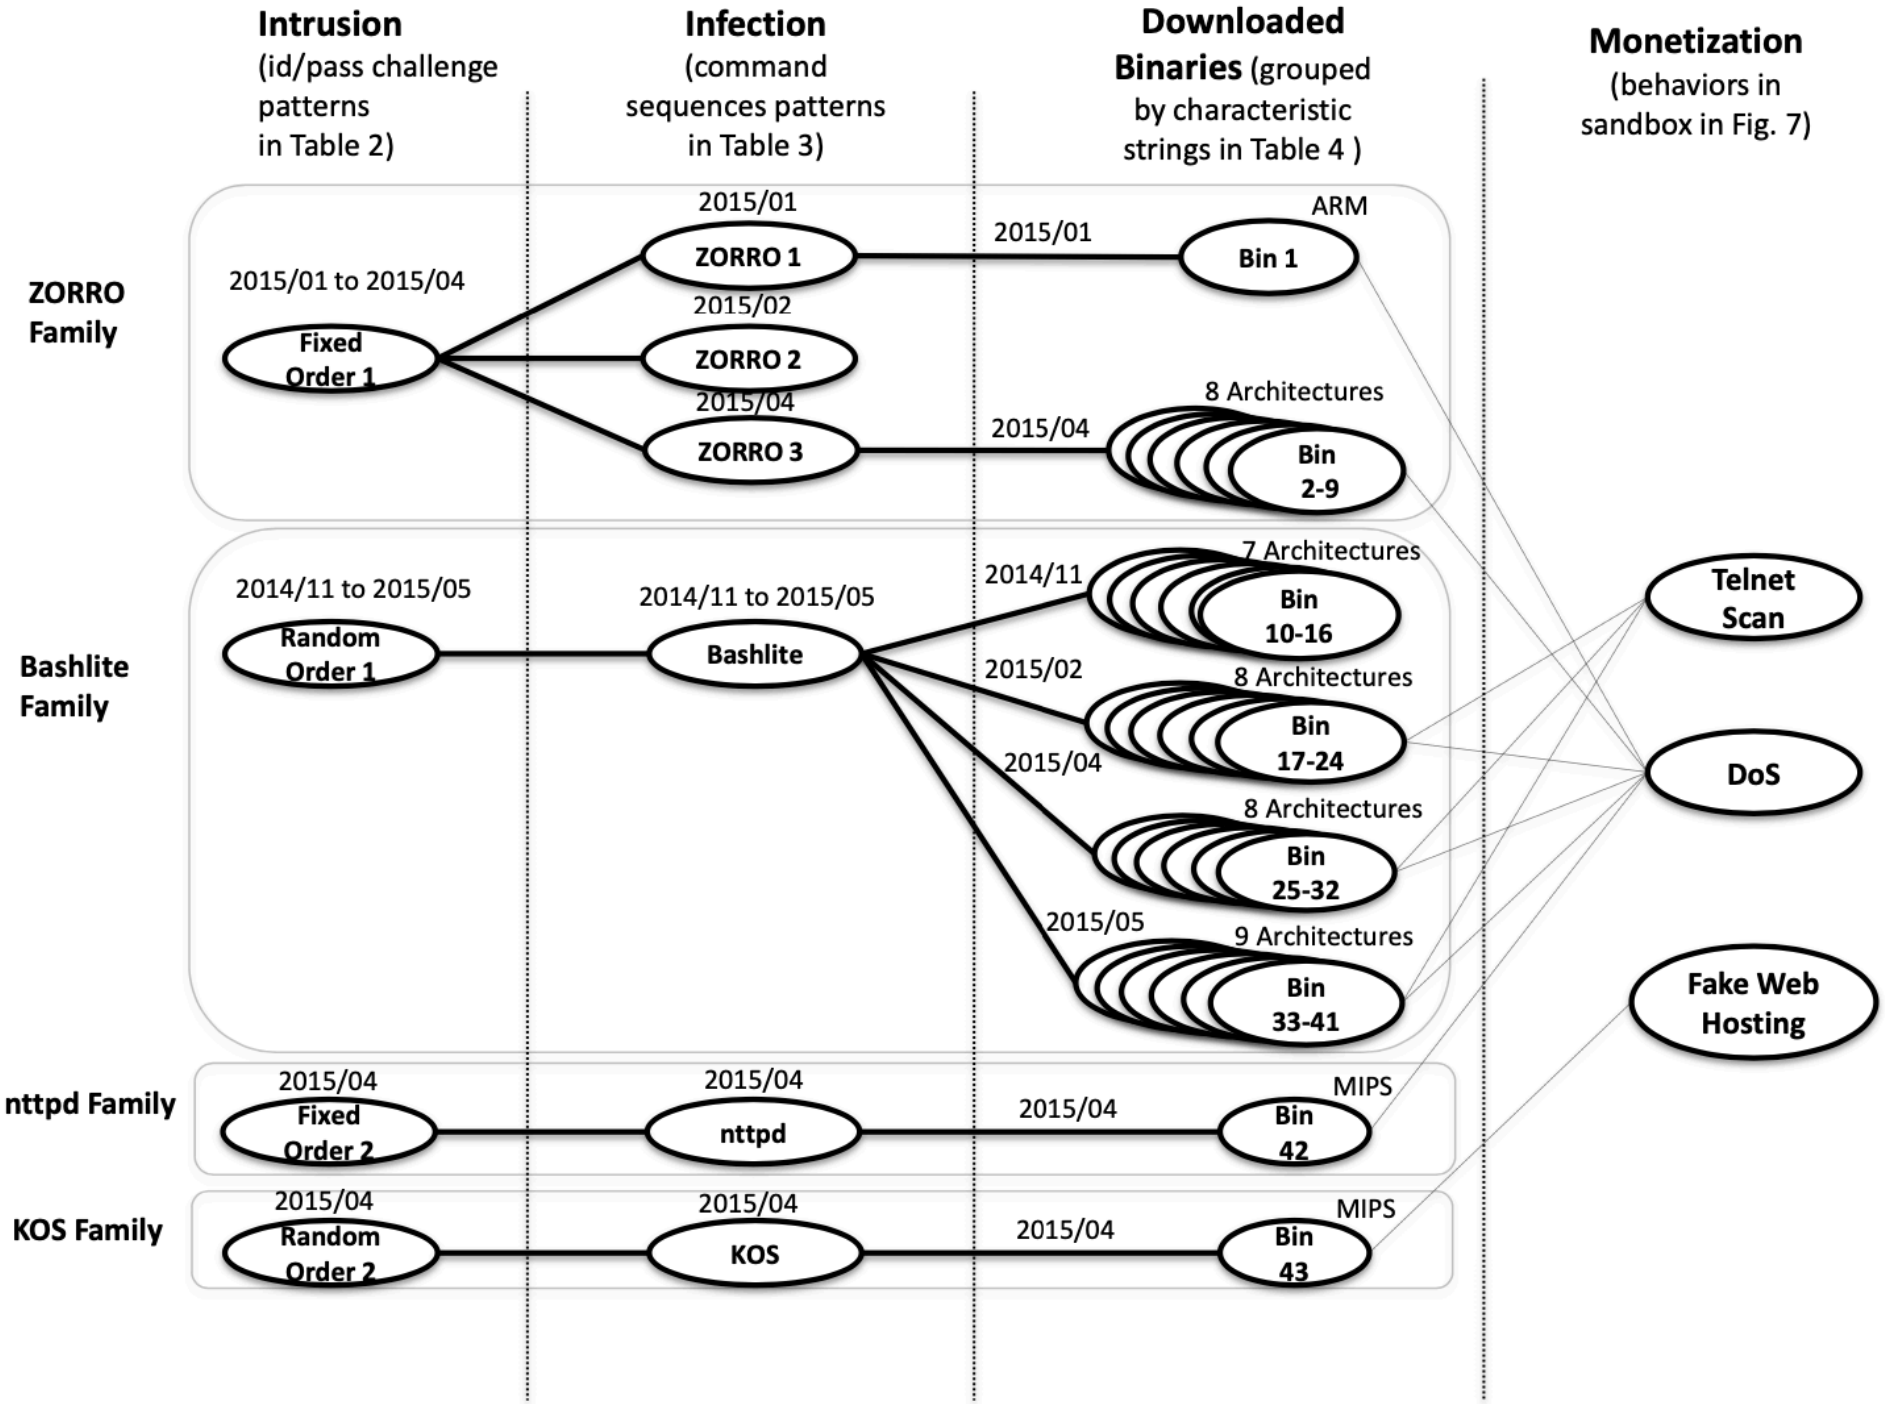
Intrusion
(id/pass challenge patterns in Table 2)
Infection
(command sequences patterns in Table 3)
Downloaded Binaries (grouped by characteristic strings in Table 4)
Monetization
(behaviors in sandbox in Fig. 7)
ZORRO Family
2015/01 to 2015/04
Fixed Order 1
2015/01
ZORRO 1
2015/02
ZORRO 2
2015/04
ZORRO 3
2015/01
ARM
Bin 1
2015/04
8 Architectures
Bin 2-9
Bashlite Family
2014/11 to 2015/05
Random Order 1
2014/11 to 2015/05
Bashlite
2014/11
7 Architectures
Bin 10-16
2015/02
8 Architectures
Bin 17-24
2015/04
8 Architectures
Bin 25-32
2015/05
9 Architectures
Bin 33-41
nttpd Family
2015/04
Fixed Order 2
2015/04
nttpd
2015/04
MIPS
Bin 42
KOS Family
2015/04
Random Order 2
2015/04
KOS
2015/04
MIPS
Bin 43
Telnet Scan
DoS
Fake Web Hosting

Results

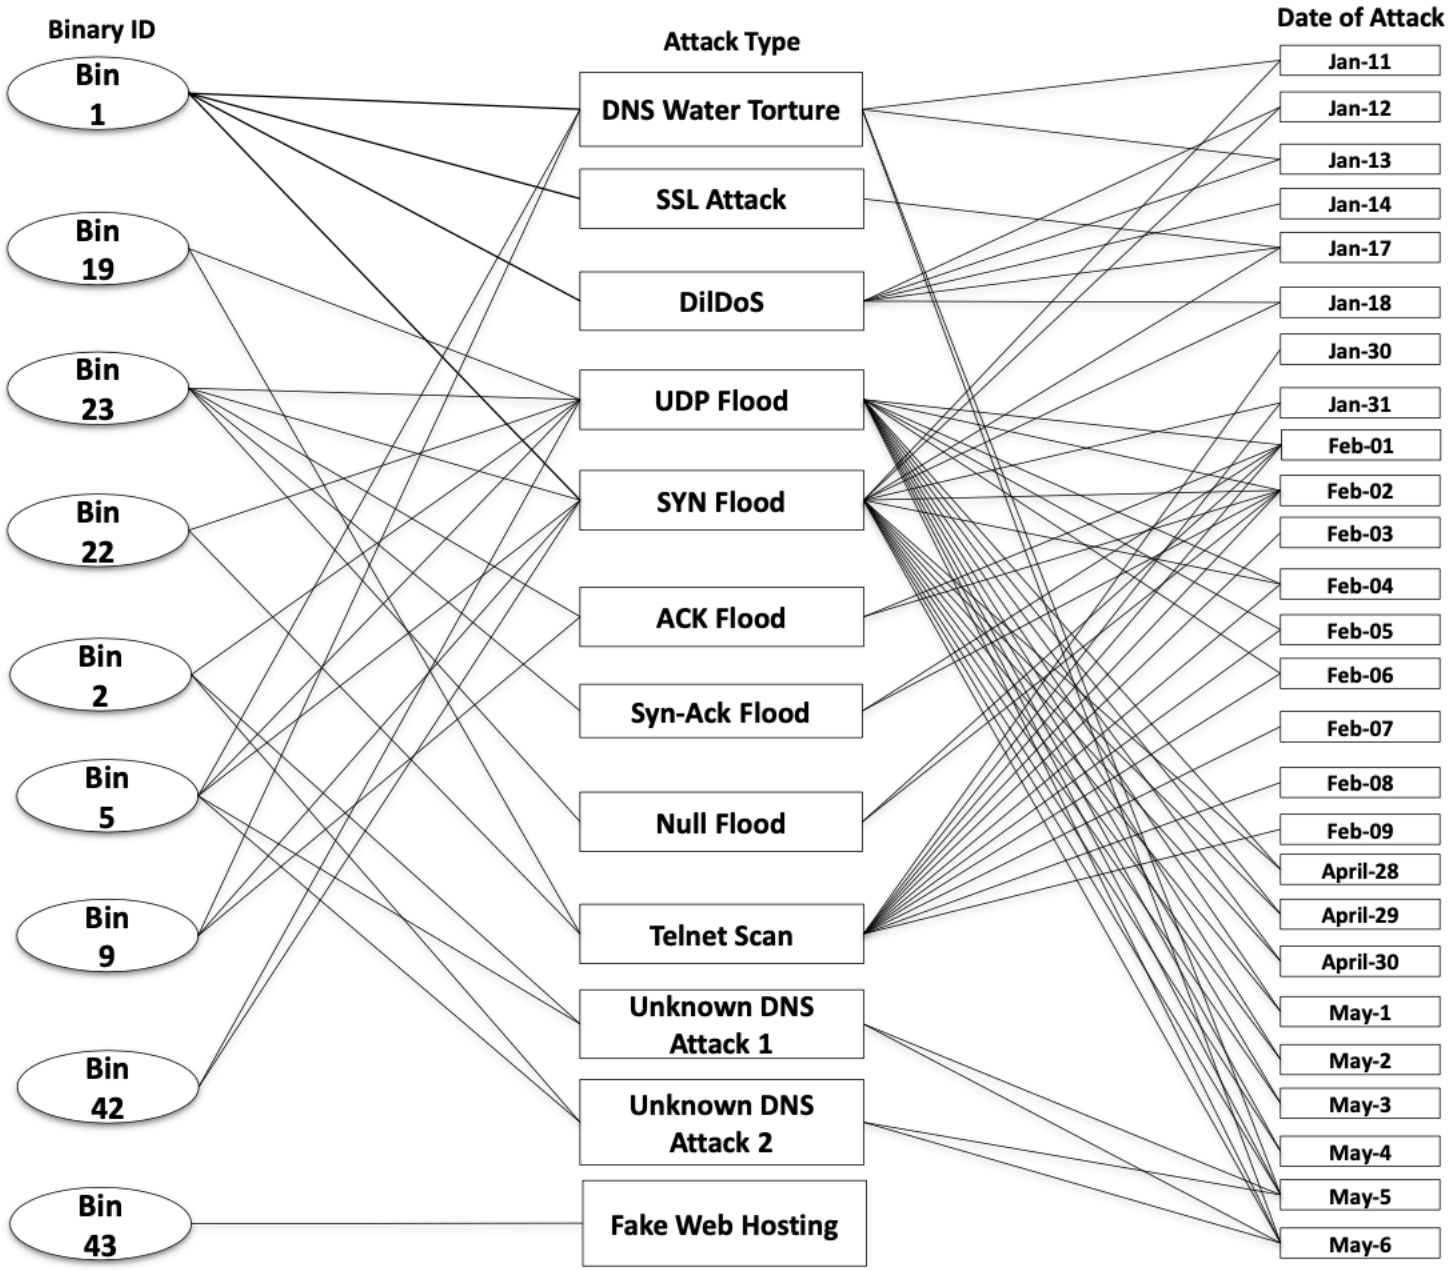

UNIVERSITY OF TWENTE.

SDN LABS

Quiz

Most important next-step

A: More CPU architectures

B: Passthrough and monitor C&C traffic

C: Standardized botnet profiles for sharing between organizations

D: Running on real (IoT) hardware

Key takeaways

IoT world heterogeneous => honeypots more complex

High-interaction needed to get useful results

Require many (!) IP addresses to catch scans

Discussion

⇒ What is IoT about IoTPOT?

⇒ Ethical considerations in running a honeypot?

⇒ How would you improve IoTPOT?

⇒ Others means to achieve the same?

Honware: A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

Vetterl, A., & Clayton, R. (2019, November). Honware: A virtual honeypot framework for capturing CPE and IoT zero days. In *Symposium on Electronic Crime Research (eCrime). IEEE.*

UNIVERSITY OF TWENTE.

SIDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- We've seen IoTPOT as a generic example, can we improve on that model?
  - Specialized honeypots can be built for known malware (leaked Mirai sourcecode)
  - But this might not capture attack traffic of unknown derivates (e.g. Yowai/Hakai)
- Malware engineers can easily scan the whole IPv4 Internet to look for vulnerable devices and quickly infect them.
- This means defenders need to scale fast too
  - IoTPOT → Hardcoded answers (and limited sandbox), Firmadyne → Not setup for network traffic, SIPHON→ physical devices
- Using original firmware as a basis for honeypots

UNIVERSITY OF TWENTE.

SIDN LABS

Quiz 1

How long does it take to scan the whole IPv4 space?

A. Around 5 minutes
B. Around 60 minutes
C. Around 1 day
D. Around 7 days

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- Using original firmware as a honeypot basis
  - Automated firmware extraction with Binwalk
  - Customizing the kernel to allow logging & emulating proprietary hardware
  - Signal interception (signals are a form of inter-process communication (IPC))
  - Module loading disabled
  - NVRAM is not available and thus has to be emulated
  - Network configuration (adding interfaces)
  - Emulation self-check (am I reachable via ping?)

UNIVERSITY OF TWENTE.

SDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

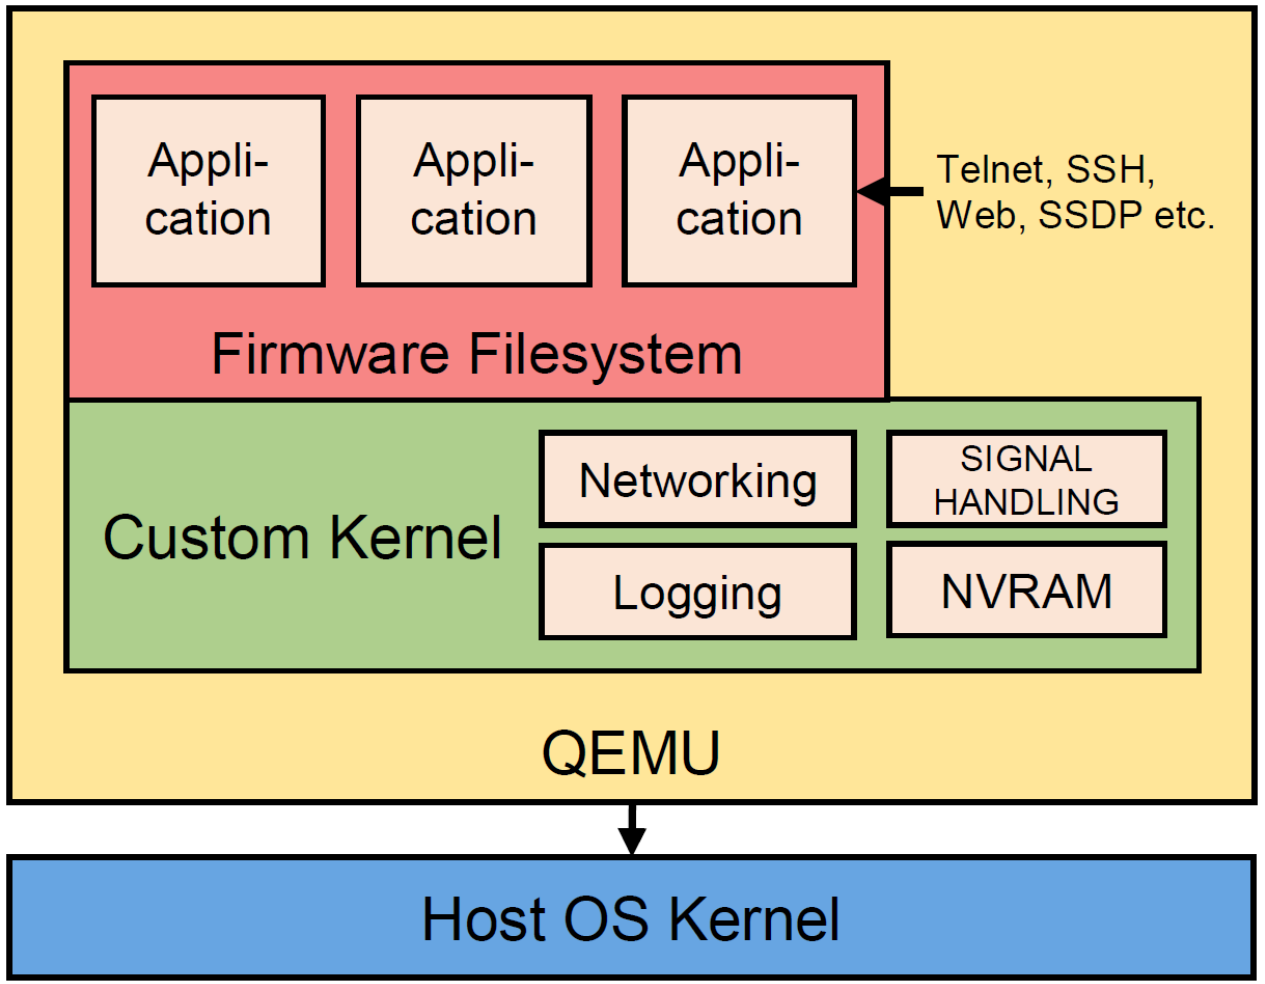

UNIVERSITY OF TWENTE.

SDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- **Not required, but fun:**
- **Reverse engineering my router's firmware with binwalk**
- https://embeddedbits.org/reverse-engineering-router-firmware-with-binwalk/
- **Playing with signals**
- **http://www.it.uu.se/education/course/homepage/os/vt18/module-2/signals/**

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- How does this system compare to the alternative (Firmadyne)?
- Out of 8387 available firmwares, 4650 could be successfully extracted (55.4%)
    - Possibly due to having weaker constraints on the size of the extracted image
- From the 4650 extracted firmware images, 1903 responded to ICMP traffic (40.9%). Firmadyne only achieved this for 460 firmware images (15.8%)
    - Likely due to the kernel customizations, and handling of crashes

UNIVERSITY OF TWENTE.

SIDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

# Brand	Available (2019-03/2016-02/Δ)		Extracted (honw./firm./Δ)		Network reach. (honw./firm./Δ)	
1 Actiontec	0/14	14↓	–	–	–	–
2 Airlink101	0/15	15↓	–	–	–	–
3 Apple	0/9	9↓	–	–	–	–
4 Asus	1/3	2↓	1/1	←	1/0	1↑
5 AT&T	3/25	22↓	0/2	2↓	–	–
6 AVM	0/132	132↓	–	–	–	–
7 Belkin	123/140	17↓	49/49	←	9/0	9↑
8 Buffalo	97/143	46↓	6/7	1↓	2/1	1↑
9 CenturyLink	13/31	18↓	7/4	3↑	7/0	7↑
10 Cerowrt	0/14	14↓	–	–	–	–
11 Cisco	0/61	61↓	–	–	–	–
12 D-Link	1443/4688	3245↓	537/498	39↑	272/115	157↑
13 Forceware	0/2	2↓	–	–	–	–
14 Foscam	44/56	12↓	5/5	←	–	–
15 Haxorware	0/7	7↓	–	–	–	–
16 Huawei	13/29	16↓	0/3	3↓	–	–
17 Inmarsat	0/47	47↓	–	–	–	–
18 Iridium	0/17	17↓	–	–	–	–
19 Linksys	32/126	94↓	26/26	←	15/1	14↑
20 MikroTik	4/13	9↓	–	–	–	–
21 Netgear	1396/5280	3884↓	639/629	10↑	384/187	197↑
22 On Networks	0/28	28↓	–	–	–	–
23 Open Wir.	0/1	1↓	–	–	–	–
24 OpenWrt	756/1498	742↓	714/705	9↑	674/0	674↑
25 pfSense	214/256	42↓	–	–	–	–
26 Polycom	612/644	32↓	0/24	24↓	–	–
27 QNAP	8/464	456↓	–	–	–	–
28 RouterTech	0/12	12↓	–	–	–	–
29 Seiki	0/16	16↓	–	–	–	–
30 Supermicro	0/150	150↓	–	–	–	–
31 Synology	1977/2094	117↓	1866/239	1627↑	–	–
32 Tenda	6/244	238↓	4/3	1↑	2/0	2↑
33 Tenvis	9/49	40↓	6/6	←	6/4	2↑
34 Thuraya	0/18	18↓	–	–	–	–
35 Tomato	362/2942	2580↓	362/362	←	217/0	217↑
36 TP-Link	463/1072	609↓	171/171	←	147/95	52↑
37 TRENDnet	336/822	486↓	134/100	34↑	87/37	50↑
38 Ubiquiti	26/51	25↓	20/19	1↑	11/0	11↑
39 u-blox	0/16	16↓	–	–	–	–
40 Verizon	0/37	37↓	–	–	–	–
41 Western Dig.	0/1	1↓	–	–	–	–
42 ZyXEL	449/1768	1319↓	103/67	36↑	69/20	49↑
Total	8387/23035	14648↓	4650/2920	1730↑	1903/460	1443↑

UNIVERSITY OF TWENTE.

SIDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

TABLE II
COMPARING HONWARE AND FIRMADYNE: TOP 15 LISTENING SERVICES.

Prot.	Port/Service	Honware	Firmadyne	Δ
TCP	23/telnet	879	149	730↑
TCP	80/http	676	293	383↑
UDP	67/dhcp	316	160	156↑
UDP	1900/UPnP	239	128	111↑
UDP	53/various	239	174	65↑
TCP	3333/dec-notes	222	102	120↑
TCP	5555/freeciv	203	57	146↑
TCP	5431/UPnP	177	48	129↑
UDP	137/netbios	154	82	72↑
TCP	53/domain	139	73	66↑
TCP	443/https	107	105	2↑
UDP	5353/mdns	102	34	68↑
UDP	69/tftp	104	26	78↑
TCP	1900/UPnP	56	60	4↓
TCP	49152/UPnP	53	62	9↓

UNIVERSITY OF TWENTE.

SIDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- How does this system compare to the real deal (hardware in the wild)?
- Fingerprinting of honeypots is an ongoing concern

UNIVERSITY OF TWENTE.

SDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

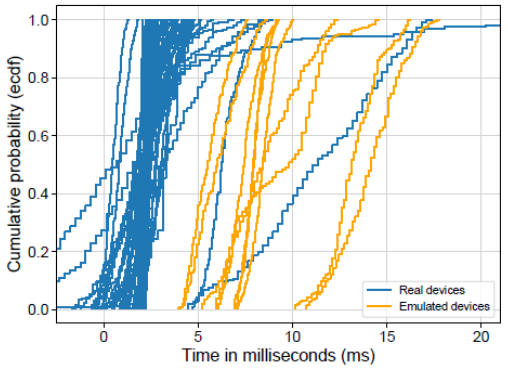

(a) ASUS RT-AC52U FTP server: Time to welcome message

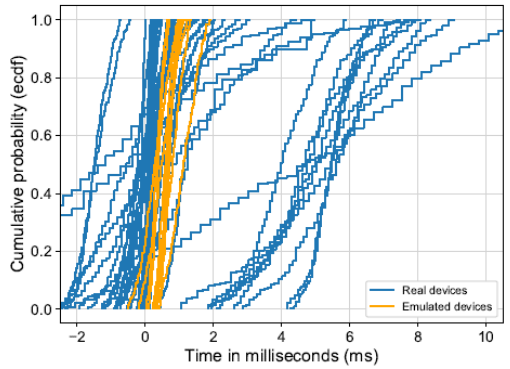

(b) ASUS RT-AC52U FTP server: Time between resource request (carriage return) and login message

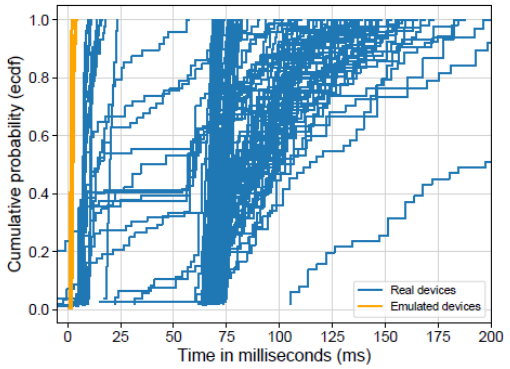

(c) Zyxel VMG1312-B10A Telnet server: Time to telnet negotiation characters

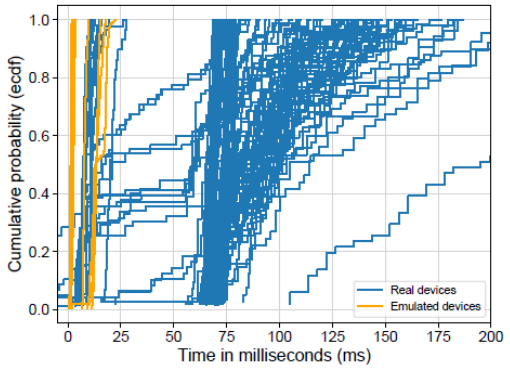

(d) Zyxel VMG1312-B10A Telnet server: Time to Login message

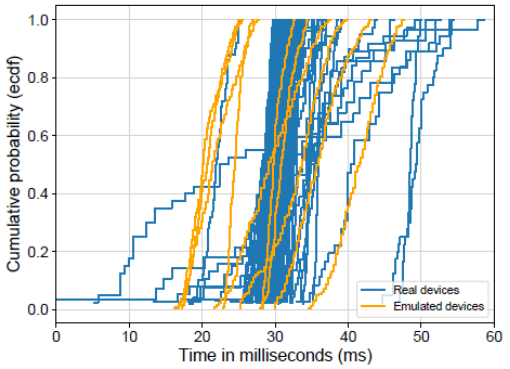

(e) D-Link DIR-825 HTTPS server: Time to complete the TLS handshake

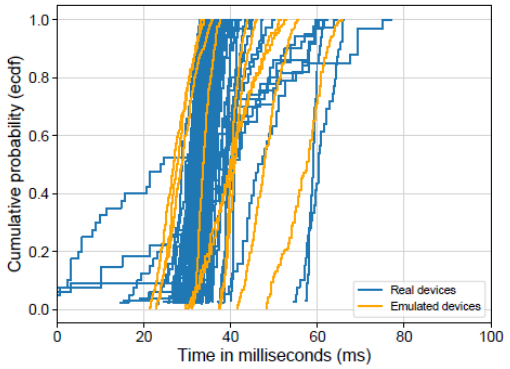

(f) D-Link DIR-825 HTTPS server: Time between ClientHello and resource received (web page)

UNIVERSITY OF TWENTE.

SDN LABS

Quiz 2

Hosting the honeypots in the cloud can aid attackers in the fingerprinting process

A. True
B. False

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- Real world results: fast

1. UPnPHunter took a research team 1 month to reverse engineer, Honware detected the complete attack within 24 hours
2. DNS hijack, a previously unknown attack
3. UPnPProxy
4. Mirai variants, target port 80 (HTTP) instead of 23 (Telnet)

- Detected malware samples were unknown to the wider community (Virustotal)

UNIVERSITY OF TWENTE.

SDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

GET /cgi-bin/timepro.cgi?tmenu=netconf&smenu=wansetup&act=save&wan=wan1&ifname=eth1&sel=dynamic&wan_type=dynamic&allow_private=on&dns_dynamic_chk=on&userid=&passwd=&mtu.pppoe.eth1=1454&lcp_flag=1&lcp_echo_interval=30&lcp_echo_failure=10&mtu.static.eth1=1500&fdns_dynamic1=185&fdns_dynamic2=117&fdns_dynamic3=74&fdns_dynamic4=100&sdns_dynamic1=185&sdns_dynamic2=117&sdns_dynamic3=74&sdns_dynamic4=101 HTTP/1.1

/sbin/iptables -t nat -A PREROUTING -i br0 -d 192.168.0.1 -p udp --dport 53 -j DNAT --to-destination 185.117.74.100

118.30.28.10
AS41718: China Great Firewall Network Limited Company

>40 IPs with the same certificate

UNIVERSITY OF TWENTE.

SDN LABS

A Virtual Honeypot Framework for Capturing CPE and IoT Zero Days

- At the beginning we were not able to capture a valid sample as the honeypot needs to be able to simulate the above scenarios. We had to tweak and customize our honeypot quite a few times, then finally in Oct, we got it right and successfully tricked the botnet to send us the sample (we call it BCMUPnP_Hunter).
- https://blog.netlab.360.com/bcmpupnp_hunter-a-100k-botnet-turns-home-routers-to-email-spammers-en/
- Original slides by the authors of the paper:
- https://www.cl.cam.ac.uk/~amv42/papers/vetterl-clayton-honware-virtual-honeypot-framework-ecrime-19-slides.pdf

UNIVERSITY OF TWENTE.

SDN LABS

Conclusion

- Honware uses real services/applications which are shipped with the device
  - In addition to that, the native configuration files are loaded
- Better than existing emulation strategies in all areas
  - Extraction, network reachability, listening services
- Capable of detecting vulnerabilities at scale
  - Rapid emulation cuts the attackers' ability to exploit vulnerabilities for considerable time

UNIVERSITY OF TWENTE.

SIDN LABS

Discussion of honeypot frameworks

1. What do you think of the proposed frameworks today? Would you change something and why?

2. Let's link this back to the lecture of governance and regulation:
   Should governments only allow the sale of an IoT device, if they can run the firmware on a testbench?

Volg ons

SIDN.nl

@SIDN

SIDN

Discussion & feedback

Next lecture: **Wed May 27, 10:45-12:30**
Topic: IoT edge security systems

UNIVERSITY OF TWENTE.

SIDN LABS