

Lecture #3: IoT Botnet Measurements

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

University of Twente | May 18, 2022

For 8 years, a hacker operated a massive IoT botnet just to download Anime videos



CITY
SITE.



Admin

Interactive lectures

- 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 questions (not graded) and discussion
 - We ask at least one of you to share their thoughts on each paper (pros, cons, surprises)
 - 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

- One summary for each of the papers we'll discuss during the lectures
- 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
- Due **before 7AM** on the **day of the lecture** in which the papers will be discussed
- Submit through Canvas



Schedule

No.	Date	Contents
1	Apr 26	Course introduction Guest lecture #1: IoT and SPIN
2	May 11	Lecture: IoT Concepts and Applications
3	May 18	Lecture: IoT Botnet Measurements
4	May 24	Guest lecture #2: Intro to cyber-physical systems (Jeroen Gaiser, Rijkswaterstaat)
5	May 25	Lecture: IoT Malware Analysis
6	Jun 1	Lecture: IoT Edge Security Systems
7	Jun 7	Lecture: IoT Device Security
8	Jun 14	Guest lecture #3: Strengthening the IoT Ecosystem: Privacy Preserving IoT Security Management (Dr Anna Maria Mandalari, Imperial College London)
9	Jun 15	Lecture: IoT in Non-Carpeted Areas
10	Jun 22	Lecture: IoT Honeypots (re-sit)

Important dates

- Two summaries per lecture: before the lecture (07:00) in which the papers will be discussed
- Lab report (PDF) and required files: **Sun June 26, 2022, 23:59 CEST**
- All to be submitted through CANVAS

Introduction to today's lecture

Today's objective

- Discussing two botnets: after the lecture, you will be able to discuss how IoT botnets are organized and spread their infections.
- [Mirai] is the infamous botnet that alerted many of the risks of IoT devices.
- [Hajime] is a more advanced IoT botnet, compared to Mirai, when it comes to bot management and usage of exploits.
- Contributes to SSI learning goal #1: “Understand IoT concepts and applications, security threats, technical solutions, and a few relevant standardization efforts in the IETF”

Today's papers

Are about measuring IoT botnets

- **[Mirai]** M. Antonakakis, T. April, M. Bailey, M. Bernhard, E. Bursztein, J. Cochran, Z. Durumeric, J. A. Halderman, L. Invernizzi, M. Kallitsis, D. Kumar, C. Lever, Z. Ma, J. Mason, D. Menscher, C. Seaman, N. Sullivan, K. Thomas, and Y. Zhou, “Understanding the Mirai Botnet”, in: 26th USENIX Security Symposium, 2017
- **[Hajime]** S. Herwig, K. Harvey, G. Hughey, R. Roberts, and D. Levin, “Measurement and Analysis of Hajime, a Peer-to-peer IoT Botnet”, Network and Distributed Systems Security (NDSS) Symposium 2019, San Diego, CA, USA, February 2019

“Understanding the Mirai Botnet” 26th USENIX Security Symposium, 2017

**Antonakakis, April, Bailey, Bernhard, Bursztein, Cochran,
Durumeric, Halderman, Invernizzi, Kallitsis, Kumar, Lever, Ma,
Mason, Menscher, Seaman, Sullivan, Thomas, and Zhou**

Mirai post-mortem

- Impressive cooperation between = different vantage points:
 - Akamai Technologies, Cloudflare, Google, Merit Network
 - Georgia Institute of Technology, University of Illinois Urbana-Champaign, University of Michigan



Quiz

Botnets can be used for purposes other than launching DDoS attacks.

For what other activity was the Mirai botnet used?

- A Bitcoin mining
- B Sending spam
- C Sharing videos
- D Click fraud

Mirai uses default passwords

```
// Set up passwords
add_auth_entry("\x50\x40\x40\x56", "\x5A\x41\x11\x17\x13\x13", 10); // root xc3511
add_auth_entry("\x50\x40\x40\x56", "\x54\x4B\x58\x5A\x54", 9); // root vizxv
add_auth_entry("\x50\x40\x40\x56", "\x43\x46\x4F\x4B\x4C", 8); // root admin
add_auth_entry("\x43\x46\x4F\x4B\x4C", "\x43\x46\x4F\x4B\x4C", 7); // admin admin
add_auth_entry("\x50\x40\x40\x56", "\x1A\x1A\x1A\x1A\x1A\x1A", 6); // root 888888
add_auth_entry("\x50\x40\x40\x56", "\x5A\x4F\x4A\x46\x4B\x52\x41", 5); // root xmhdipc
add_auth_entry("\x50\x40\x40\x56", "\x46\x47\x44\x43\x57\x4E\x56", 5); // root default
add_auth_entry("\x50\x40\x40\x56", "\x48\x57\x43\x4C\x56\x47\x41\x4A", 5); // root juantech
add_auth_entry("\x50\x40\x40\x56", "\x13\x10\x11\x16\x17\x14", 5); // root 123456
add_auth_entry("\x50\x40\x40\x56", "\x17\x16\x11\x10\x13", 5); // root 54321
add_auth_entry("\x51\x57\x52\x52\x40\x50\x56", "\x51\x57\x52\x52\x40\x50\x56", 5); // support support
add_auth_entry("\x50\x40\x40\x56", "", 4); // root (none)
add_auth_entry("\x43\x46\x4F\x4B\x4C", "\x52\x43\x51\x51\x55\x40\x50\x46", 4); // admin password
add_auth_entry("\x50\x40\x40\x56", "\x50\x40\x40\x56", 4); // root root
add_auth_entry("\x50\x40\x40\x56", "\x13\x10\x11\x16\x17", 4); // root 12345
add_auth_entry("\x57\x51\x47\x50", "\x57\x51\x47\x50", 3); // user user
add_auth_entry("\x43\x46\x4F\x4B\x4C", "", 3); // admin (none)
add_auth_entry("\x50\x40\x40\x56", "\x52\x43\x51\x51", 3); // root pass
add_auth_entry("\x43\x46\x4F\x4B\x4C", "\x43\x46\x4F\x4B\x4C\x13\x10\x11\x16", 3); // admin admin1234
```

Scanning the Internet

```
while (o1 == 127 ||
(o1 == 0) ||
(o1 == 3) ||
(o1 == 15 || o1 == 16) ||
(o1 == 56) ||
(o1 == 10) ||
(o1 == 192 && o2 == 168) ||
(o1 == 172 && o2 >= 16 && o2 < 32) ||
(o1 == 100 && o2 >= 64 && o2 < 127) ||
(o1 == 169 && o2 > 254) ||
(o1 == 198 && o2 >= 18 && o2 < 20) ||
(o1 >= 224) ||
(o1 == 6 || o1 == 7 || o1 == 11 ||
o1 == 21 || o1 == 22 || o1 == 26 ||
o1 == 28 || o1 == 29 || o1 == 30 ||
o1 == 33 || o1 == 55 || o1 == 214 ||
o1 == 215)
);
```

// 127.0.0.0/8 - Loopback
// 0.0.0.0/8 - Invalid address space
// 3.0.0.0/8 - General Electric Company
// 15.0.0.0/7 - Hewlett-Packard Company
// 56.0.0.0/8 - US Postal Service
// 10.0.0.0/8 - Internal network
// 192.168.0.0/16 - Internal network
// 172.16.0.0/14 - Internal network
// 100.64.0.0/10 - IANA NAT reserved
// 169.254.0.0/16 - IANA NAT reserved
// 198.18.0.0/15 - IANA Special use
// 224.*.*.* - Multicast

// Department of Defense

Scanning the Internet (2)

```
for (i = 0; i < SCANNER_RAW_PPS; i++)
{
    struct sockaddr_in paddr = {0};
    struct iphdr *iph = (struct iphdr *)scanner_rawpkt;
    struct tcphdr *tcph = (struct tcphdr *)(iph + 1);

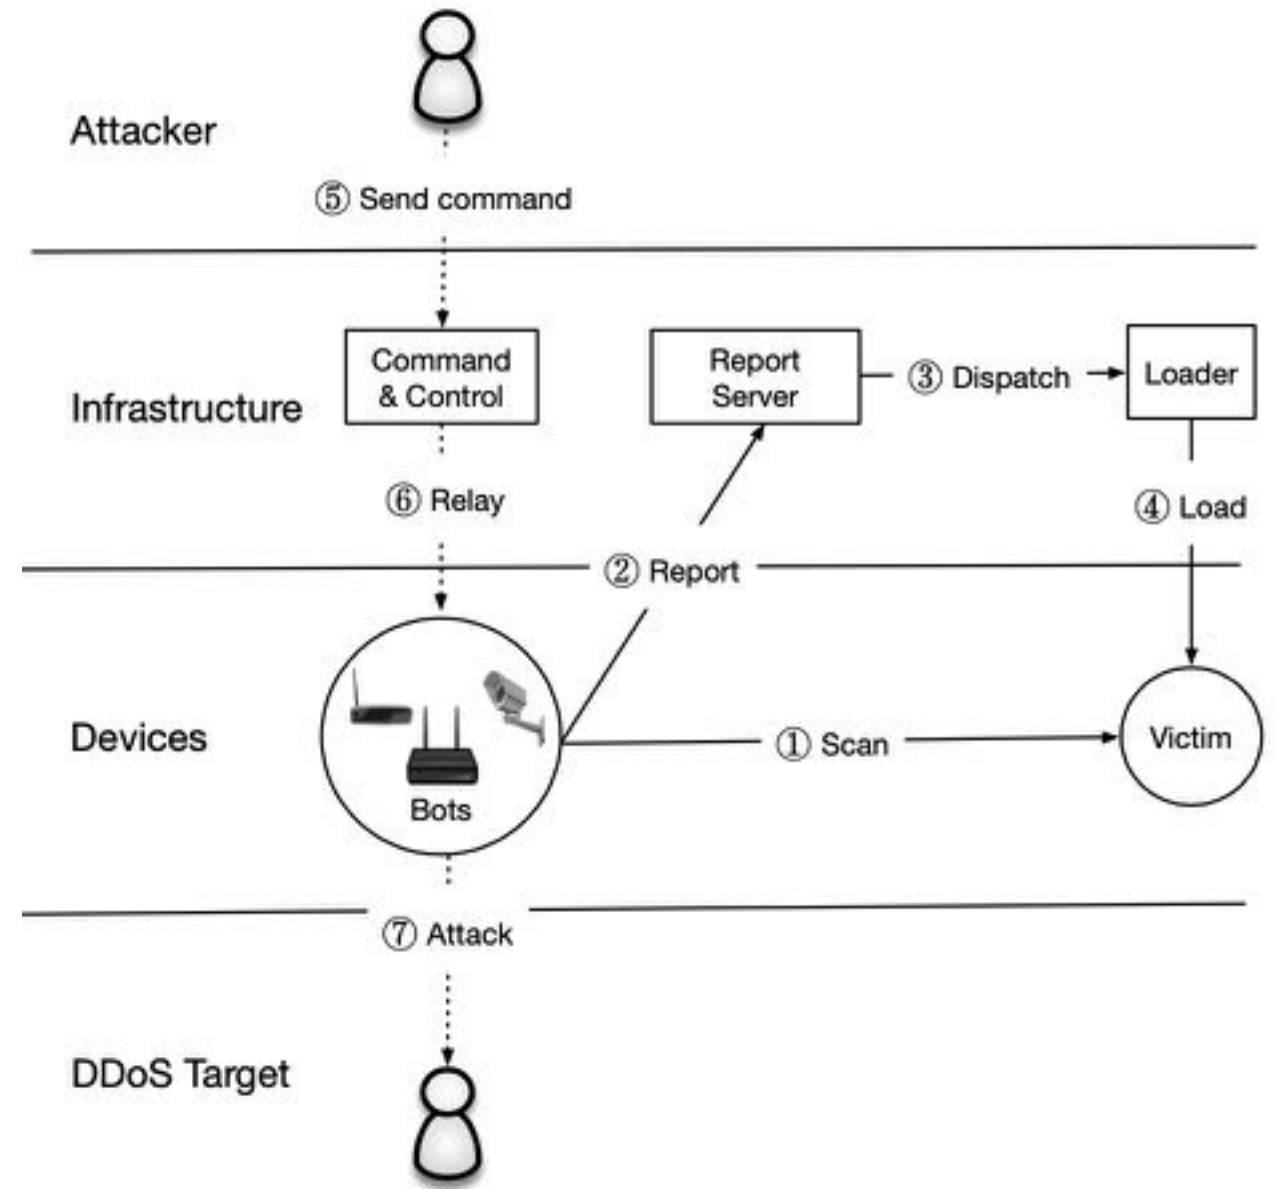
    iph->id = rand_next();
    iph->saddr = LOCAL_ADDR;
    iph->daddr = get_random_ip();
    iph->check = 0;
    iph->check = checksum_generic((uint16_t *)iph, sizeof (struct iphdr));

    if (i % 10 == 0)
    {
        tcph->dest = htons(2323);
    }
    else
    {
        tcph->dest = htons(23);
    }
    tcph->seq = iph->daddr;
    tcph->check = 0;
    tcph->check = checksum_tcpudp(iph, tcph, htons(sizeof (struct tcphdr)), sizeof (struct tcphdr));

    paddr.sin_family = AF_INET;
    paddr.sin_addr.s_addr = iph->daddr;
    paddr.sin_port = tcph->dest;
}
```


Mirai inner working

- Rapid stateless scanning: 23 and 2323 TCP SYN (seq num)
- On connection: start brute force login (10 attempts)
- Report successful login to hard-coded report server
- (Async) infect with loader program.
- Close ports and perform AV cleanup
- C2 await commands



Mirai from a network perspective

- Active scanning: (Censys)
- IoT Honeypot: 1028 unique samples and 67 C2 domains
- Passive and Active DNS to find more C2 servers
- C2 milker: 15.000 attacks




Question

What are the challenges to analyze and/or mitigate Mirai attacks?



Mirai DDoS attacks

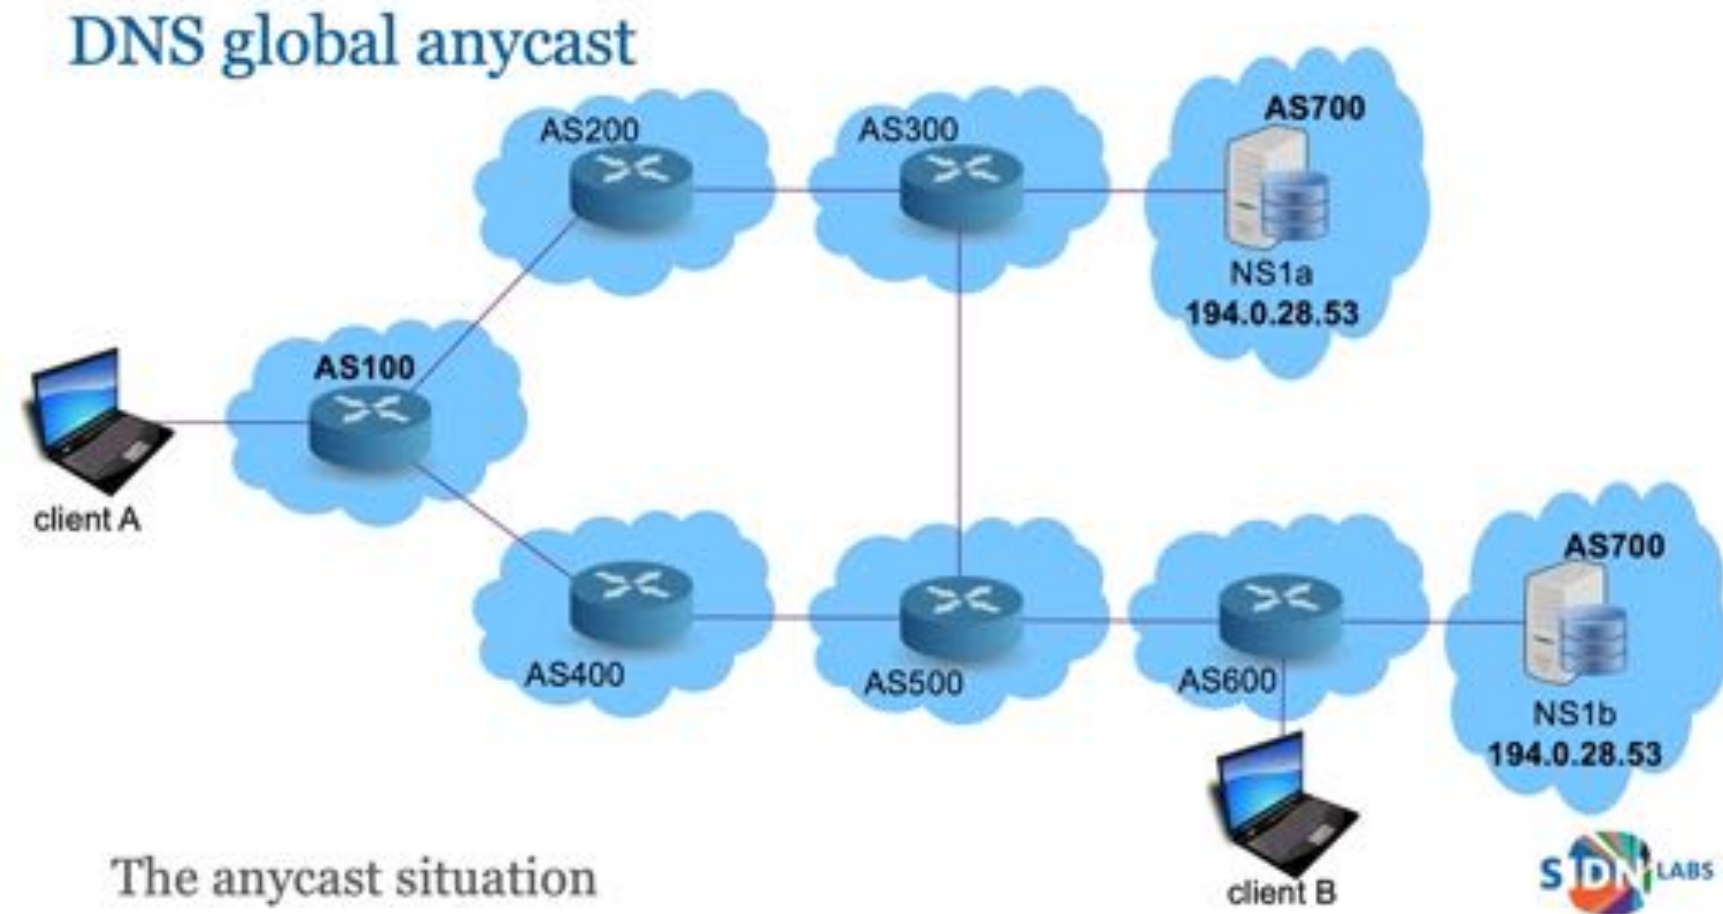
- Volumetric, TCP State Exhaustion, Application-level attacks.
- Most targets in USA (50%), France, UK.
- Games
- Mirai C2 servers
- High-profile targets: Krebs on Security, Lonestar Cell (Liberia), Dyn.



Mitigation of DDoS attacks

DDoS scrubbing service

DNS (Dyn): anycast



Lessons learned

Simple attack, lots of damage

Automatic updates

Device identification on network

IoT end-of-life devices (externality)

Connecting datasets gives a lot of information!



Question

What was the biggest 'contribution' of Mirai in your opinion?

Volg ons

 SIDN.nl

 @SIDN

 SIDN

Discussion

Measurement and Analysis of Hajime, a Peer-to-peer IoT Botnet

S. Herwig, K. Harvey, G. Hughey, R. Roberts, and D. Levin, “Measurement and Analysis of Hajime, a Peer-to-peer IoT Botnet”, Network and Distributed Systems Security (NDSS) Symposium 2019, San Diego, CA, USA, February 2019

Focus

- The important differences between Mirai and Hajime
- Backscatter data from a root DNS server
- Discussions

The 3 big differences

- Peer-to-Peer instead of centralized command & control
 - More exploits based on the Vault7 leak
 - Custom protocol to spread the malware
-
- No malicious activity had been recorded. Does this count as difference?

Architecture	Port	Service	Method
mipseb	23, 5358	Telnet	credentials
	7547	TR-064	CVE-2016-10372
	many	HTTP	Chimay-Red
	80	HTTP	CVE-2018-10561,-10562
mipsel	23, 5358	Telnet	credentials
	7547	TR-064	CVE-2016-10372
arm7	23, 5358	Telnet	credentials
	81	HTTP	GoAhead-Webs credentials
	81	HTTP	Cross Web Server RCE
arm6	23,5358	Telnet	credentials
arm5	23, 5358	Telnet	credentials
	9000	MCTP	CVE-2015-4464

TABLE I: Hajime's architecture-specific access methods and the corresponding ports scanned

P2P Mechanisms

- DHT (Kademlia) based.
 - Known from e.g. BitTorrent
 - Traditional BitTorrent connections relied on trackers to exchange seeder/leecher information
- Basically, a distributed Key-Value storage
 - Key is filename concatenated with current day' timestamp
 - Values are IPs which are infected with Hajime and allow for payload downloads

Question

- Since the key is computed based on the current day's timestamp, and bots may have incorrectly synchronized clocks, we look up keys for a five-day range (two days in the past through two days in the future).
- Do you think that this range will catch all devices?

Malicious activity(?)

- On infection, Hajime closes at least the following ports: 23 (Telnet), 5358 (WSDAPI), 5555 (Oracle Web Center Content/Freeciv), and 7547(CWMP)
- Do you remember which port/service was used by Mirai to infect devices?
- Small discussion: What do you think of the motive of the Hajime-bot author?

Custom uTorrent Transport Protocol

- Mirai was enumerable/detectable due to its custom TCP sequence field
- Hajime uses unique cryptographic public keys to allow for a count of infected hosts
- Some churn expected due to recreation of the public key, during updates to the .i module
- Still a stronger identifier, compared to weak identifiers such as IPs (ie. due to carrier grade NAT)

DNS backscatter data

- Based on trying to inject shell-commands into a NTP configuration file
- Vulnerable devices won't sanitize the input and then execute the commands, infecting the device.
- Remember how DNS lookups work? Invalid queries will be sent to the root DNS servers
 - Conveniently the researchers of the paper operate one of the root DNS servers

Question

- Do you think that Hajime is still active?
 - A. Yes
 - B. No

Demo

1. UTC timestamp
2. payload name
3. date used as input for computing the payload's DHT hash ID
4. payload DHT ID (the hash we lookup or announce on the DHT)
5. "seeder" or "leecher" (are we collecting seeders or leechers, respectively)
6. IPv4 address of seeder/leecher bot
7. port number of seeder/leecher bot

Demo (Backup)

```
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 98.43.129.55 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 109.148.173.191 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 79.161.52.82 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 24.88.23.242 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 69.112.168.236 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 108.173.178.204 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 71.210.33.221 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 24.115.107.208 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 176.14.243.44 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 71.190.197.164 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 70.119.82.44 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 184.83.113.35 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 137.25.255.15 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 185.108.162.49 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 176.110.136.21 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 62.46.102.115 62289
1620769710 atk.mipseb.1506215619 2021-05-09 1173332a85f47e1a40b15f3d77a550ff342442c2 seeder 67.251.129.160 62289
1620769711 1173332a85f47e1a40b15f3d77a550ff342442c2 5.139.3.14:49978 117710404a4f6e018508fce5f2855ef7b4b63620 115.
#1620769711 1173332a85f47e1a40b15f3d77a550ff342442c2 117710404a4f6e018508fce5f2855ef7b4b63620 5.139.3.14:49978 Tot
1620769711 1173332a85f47e1a40b15f3d77a550ff342442c2 144.91.111.37:9613 1173332a85f47e1a40b15f3d77a550fa00c24bc9 16
1620769711 1173332a85f47e1a40b15f3d77a550ff342442c2 144.91.111.37:9613 1173332a85f47e1a40b15f3d77a550fa00c24bc9 47
```

Demo (Backup)

```
1620769773 .i_armv6l.1509400182 2021-05-11 2f9f80b52e7df032562bfdc6174733006fb978e9 seeder 173.46.242.130 62289
1620769773 .i_armv6l.1509400182 2021-05-11 2f9f80b52e7df032562bfdc6174733006fb978e9 seeder 174.20.138.204 62289
1620769773 .i_armv6l.1509400182 2021-05-11 2f9f80b52e7df032562bfdc6174733006fb978e9 seeder 79.136.72.19 62289
1620769773 .i_armv6l.1509400182 2021-05-11 2f9f80b52e7df032562bfdc6174733006fb978e9 seeder 87.81.93.7 62289
1620769773 .i_armv6l.1509400182 2021-05-11 2f9f80b52e7df032562bfdc6174733006fb978e9 seeder 154.45.216.220 62289
1620769774 2f9f80b52e7df032562bfdc6174733006fb978e9 144.91.111.37:9613 2f9f80b52e7df032562bfdc6174733fa00c24bc9 173.46.242.130 62289
1620769774 2f9f80b52e7df032562bfdc6174733006fb978e9 144.91.111.37:9613 2f9f80b52e7df032562bfdc6174733fa00c24bc9 174.20.138.204 62289
1620769774 2f9f80b52e7df032562bfdc6174733006fb978e9 144.91.111.37:9613 2f9f80b52e7df032562bfdc6174733fa00c24bc9 79.136.72.19 62289
1620769774 2f9f80b52e7df032562bfdc6174733006fb978e9 144.91.111.37:9613 2f9f80b52e7df032562bfdc6174733fa00c24bc9 87.81.93.7 62289
1620769774 2f9f80b52e7df032562bfdc6174733006fb978e9 144.91.111.37:9613 2f9f80b52e7df032562bfdc6174733fa00c24bc9 154.45.216.220 62289
#1620769774 2f9f80b52e7df032562bfdc6174733006fb978e9 2f9f80b52e7df032562bfdc6174733fa00c24bc9 144.91.111.37:9613 Total seeders: 5 new
1620769785 .i_mipseb.1524631409 2021-05-11 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 seeder 174.17.14.156 62289
1620769785 .i_mipseb.1524631409 2021-05-11 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 seeder 64.121.214.41 62289
1620769785 .i_mipseb.1524631409 2021-05-11 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 seeder 79.136.72.19 62289
1620769785 .i_mipseb.1524631409 2021-05-11 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 seeder 188.17.175.246 62289
1620769785 .i_mipseb.1524631409 2021-05-11 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 seeder 81.217.115.184 62289
1620769786 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 144.91.111.37:9613 5b28b53bc218c21fdd3ec9b11e696cfa00c24bc9 174.17.14.156 62289
1620769786 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 144.91.111.37:9613 5b28b53bc218c21fdd3ec9b11e696cfa00c24bc9 64.121.214.41 62289
1620769786 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 144.91.111.37:9613 5b28b53bc218c21fdd3ec9b11e696cfa00c24bc9 79.136.72.19 62289
1620769786 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 144.91.111.37:9613 5b28b53bc218c21fdd3ec9b11e696cfa00c24bc9 188.17.175.246 62289
1620769786 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 144.91.111.37:9613 5b28b53bc218c21fdd3ec9b11e696cfa00c24bc9 81.217.115.184 62289
#1620769786 5b28b53bc218c21fdd3ec9b11e696c137a7420f8 5b28b53bc218c21fdd3ec9b11e696cfa00c24bc9 144.91.111.37:9613 Total seeders: 5 new
1620769787 .i_mipseb.1522574420 2021-05-12 09f5299a344afa50742c3f29ac7d9a163fc04b04 seeder 5.146.192.252 62289
1620769788 09f5299a344afa50742c3f29ac7d9a163fc04b04 144.91.111.37:9613 09f5299a344afa50742c3f29ac7d9afa00c24bc9 5.146.192.252 62289
#1620769788 09f5299a344afa50742c3f29ac7d9a163fc04b04 09f5299a344afa50742c3f29ac7d9afa00c24bc9 144.91.111.37:9613 Total seeders: 1 new
```

Lessons learned

1. Command-And-Control impossible to take down, without also affecting legitimate users
2. Multiple identifiers can help in mapping the extent of a botnet (uTP keys, backscatter data)
3. Abandoned botnets float through the Internet, like satellite debris around earth's orbit
4. (By proxy), manufacturers treat their security division poorly

Discussion: Botnet

- Why would the cleanup of IoT botnets take longer than for traditional bots?

Key takeaways

- Analyzing botnets properly requires many vantage points and datasets.
- Mirai ‘shook the world’ and showed potential of IoT botnets in terms of DDoS attacks.
- By leveraging an established decentralized communication protocol for command & control, Hajime circumvents traditional take-down measures for botnets.

Lecture feedback

<https://www.wooclap.com/SSILECTURE3>

1 & 2: wooclap



3: Open feedback



To what extent do you think you'll be able to discuss the reasons behind...



Click on the projected screen to start the question





Volg ons

 SIDN.nl

 @SIDN

 SIDN

Discussion & feedback

Next lecture: **Tue May 24, 1:45-12:30 CR 3B**