CSE227 – Graduate Computer Security

DDoS & Botnets

UC San Diego

Housekeeping

General course things to know

- Midpoint check-in document is due **2/14 at 11:59pm PT**
  - Introduction (frame the problem)
  - Related work section (should include ~5 – 10 relevant papers)
  - Research plan, current status, what's left to do

Today's lecture

Learning Objectives

- Learn about DDoS, botnets, and mechanisms for detecting DDoS traffic on the Internet
- Discuss the "Inferring Internet DoS Activity" paper
- Discuss the "Mirai Botnet" paper

Preliminaries

What is a Denial-of-Service attack?

What is a Denial-of-Service attack?

DoS: An attack that consumes the resources of a remote host of network, making it unavailable for normal use

What is a Distributed Denial-of-Service attack?

What is a Distributed Denial-of-Service attack?

DDoS: DoS, but distributed across many different attacking machines making it impossible to block solely based on IP address

Two types of DoS

- What is a *logic-based* DoS attack?

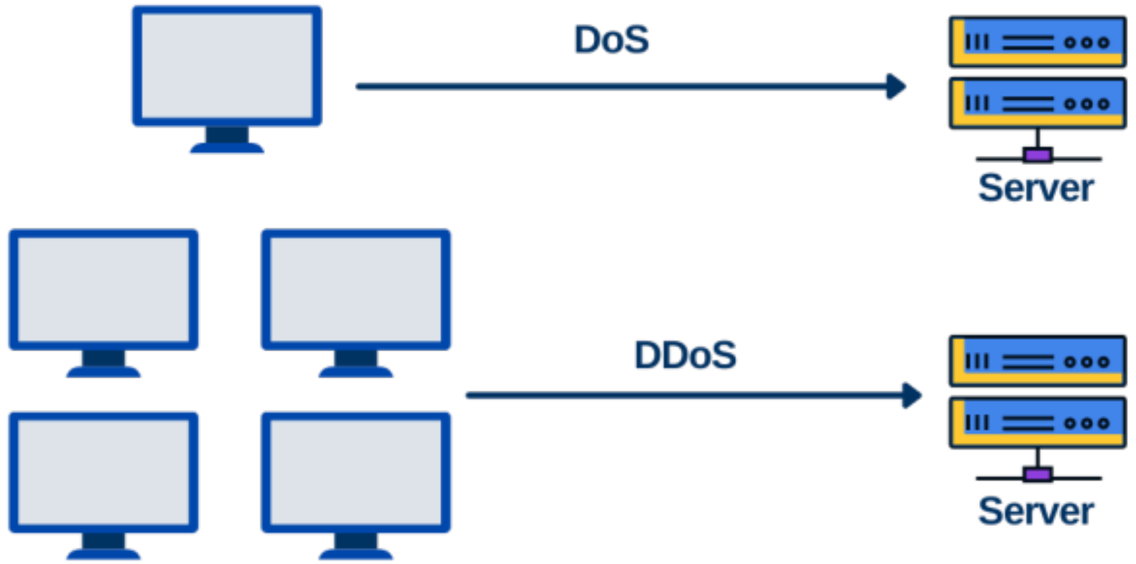

Two types of DoS

- What is a *logic-based* DoS attack?
  - Exploits some fundamental problem in the software that renders the server useless

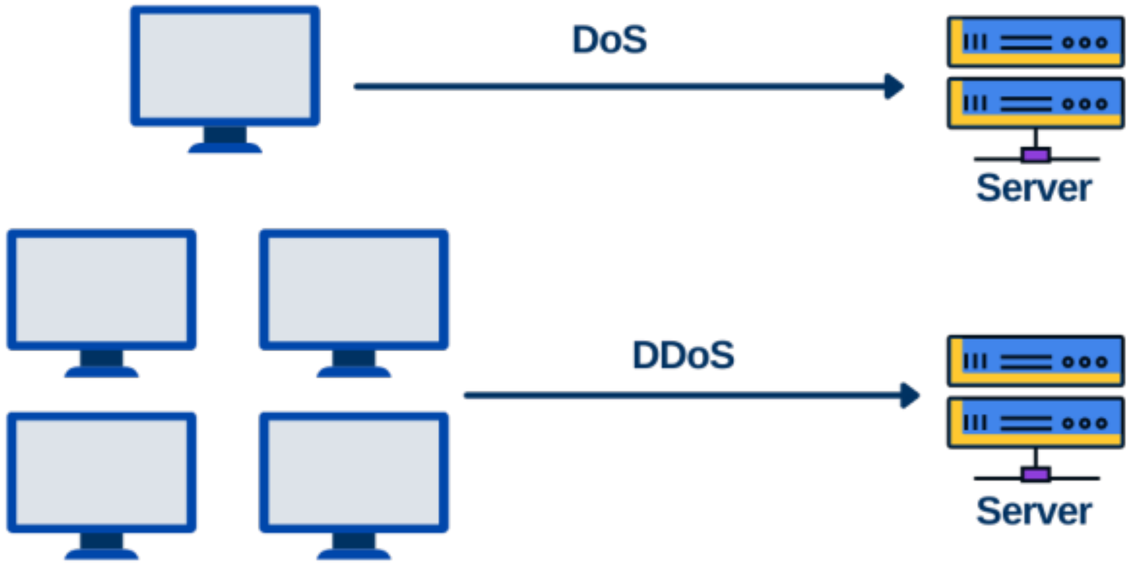

Two types of DoS

- What is a *logic-based* DoS attack?
  - Exploits some fundamental problem in the software that renders the server useless
- What is a *flooding*-based DoS attack?

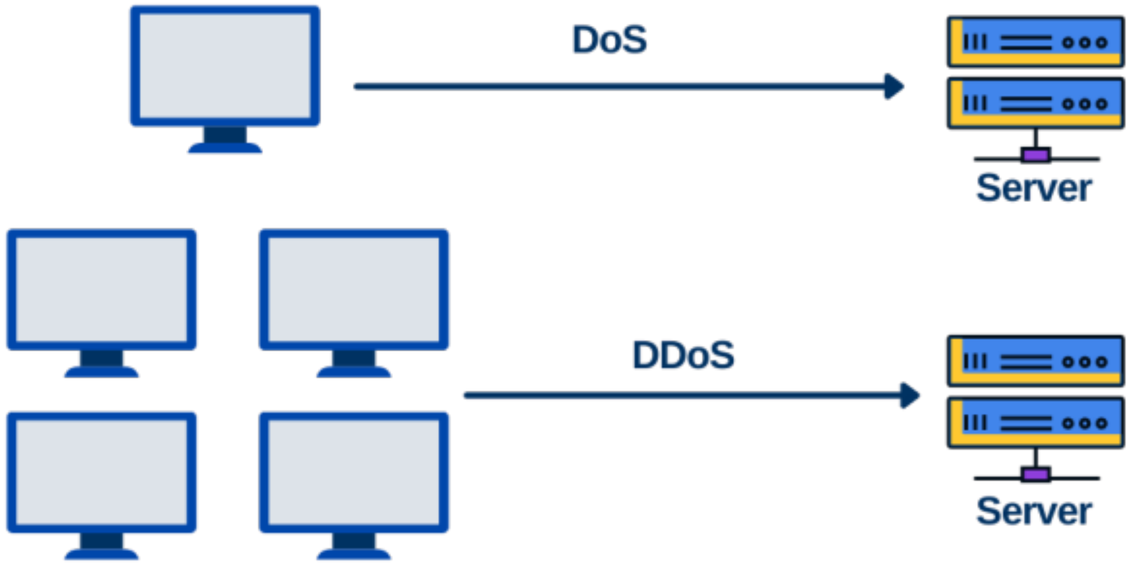

Two types of DoS

- What is a *logic-based* DoS attack?
  - Exploits some fundamental problem in the software that renders the server useless
- What is a *flooding*-based DoS attack?
  - Overwhelm resources by sending lots of packets
  - These papers: **flooding**

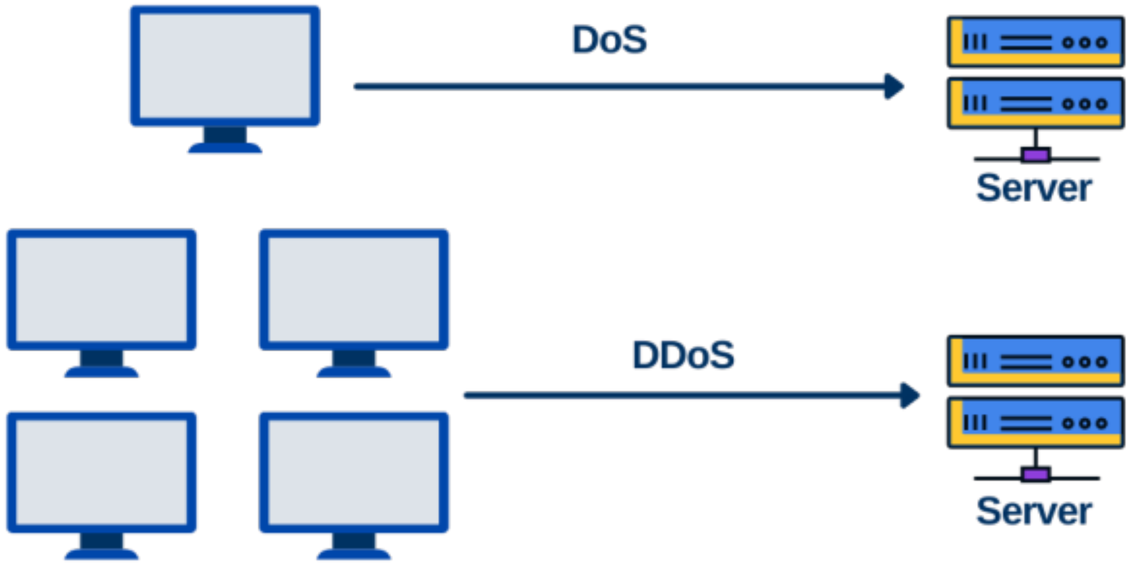

Basic Flavor of DDoS

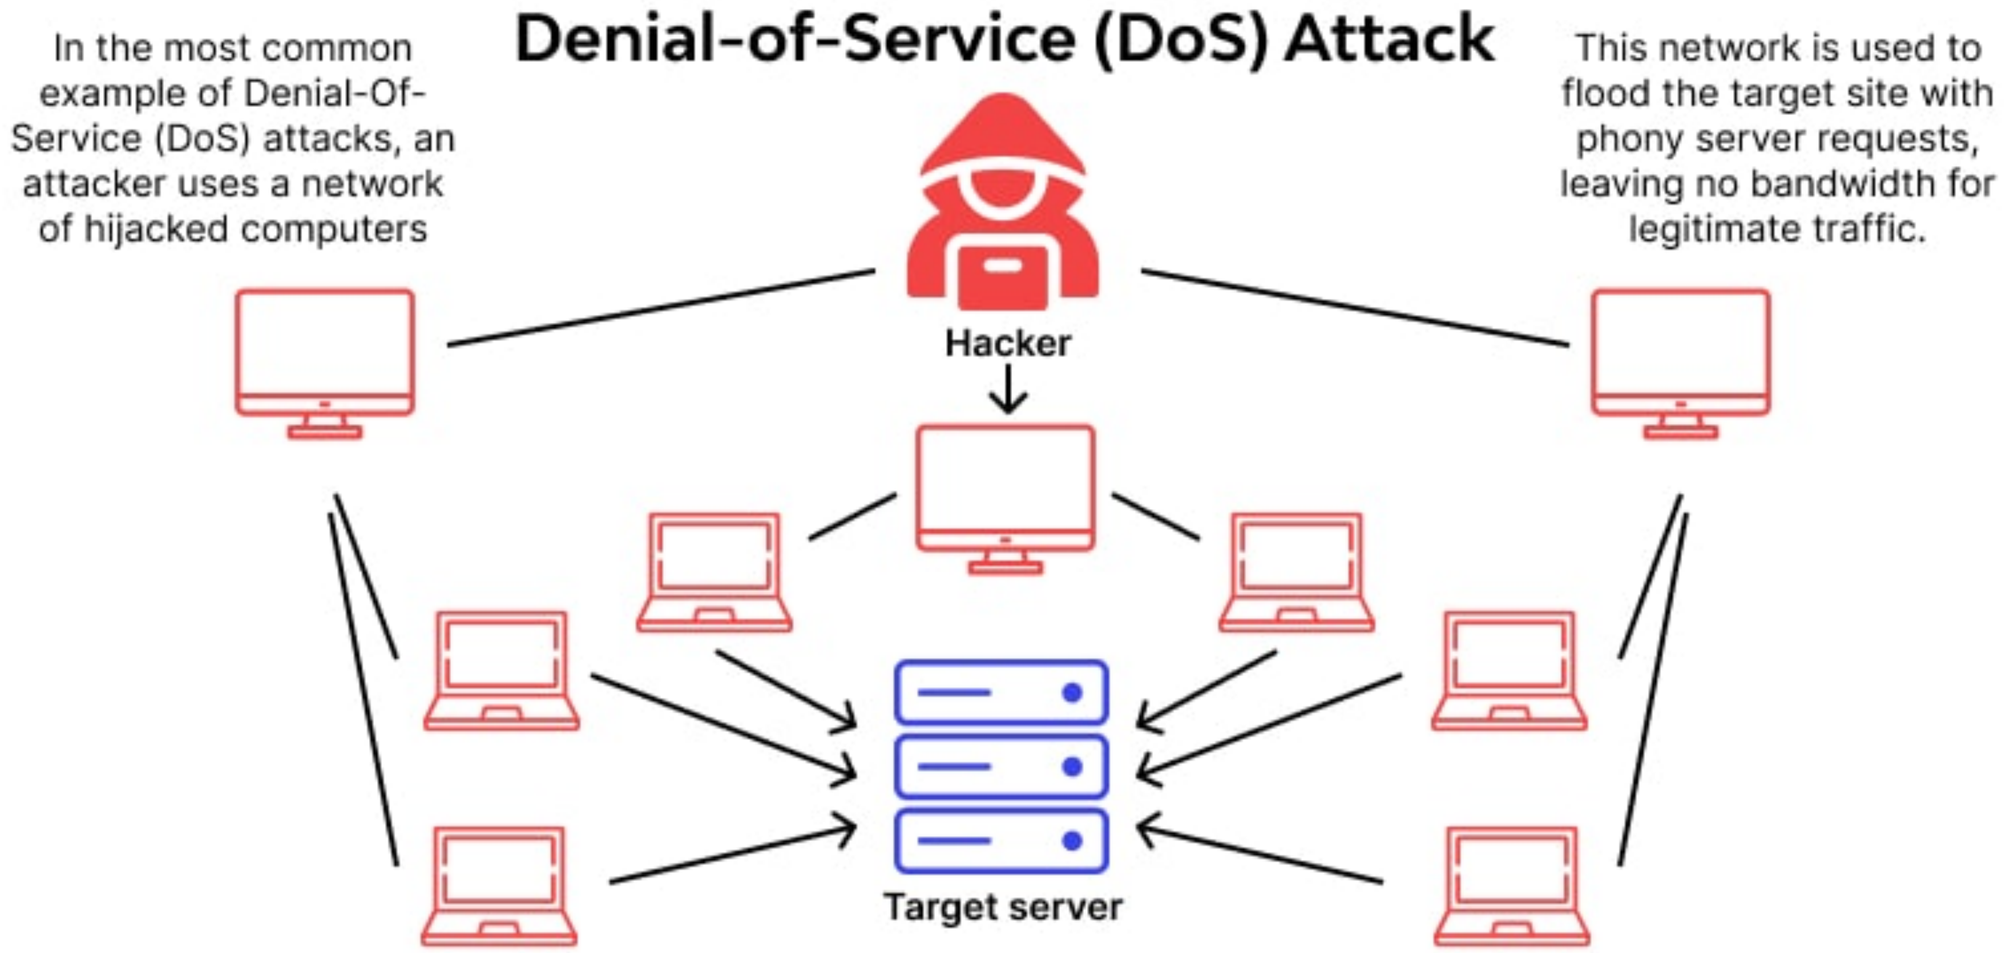

https://www.wallarm.com/what/dos-denial-of-service-attack

What is IP spoofing?

What is IP spoofing?

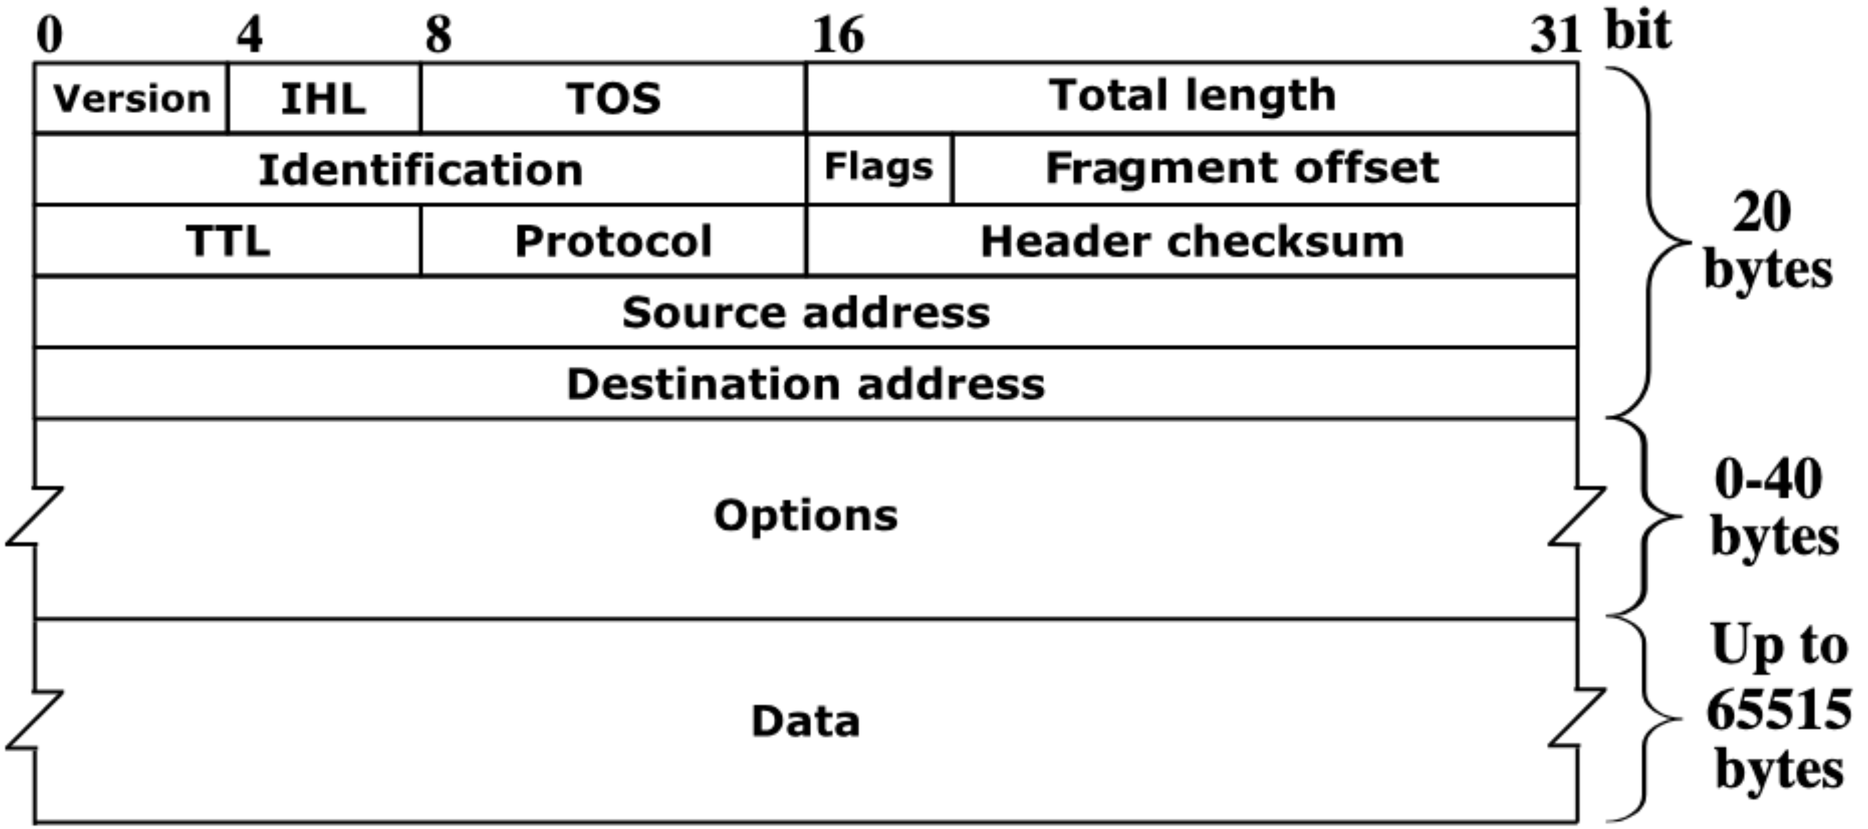

What is IP spoofing?

What is IP spoofing?

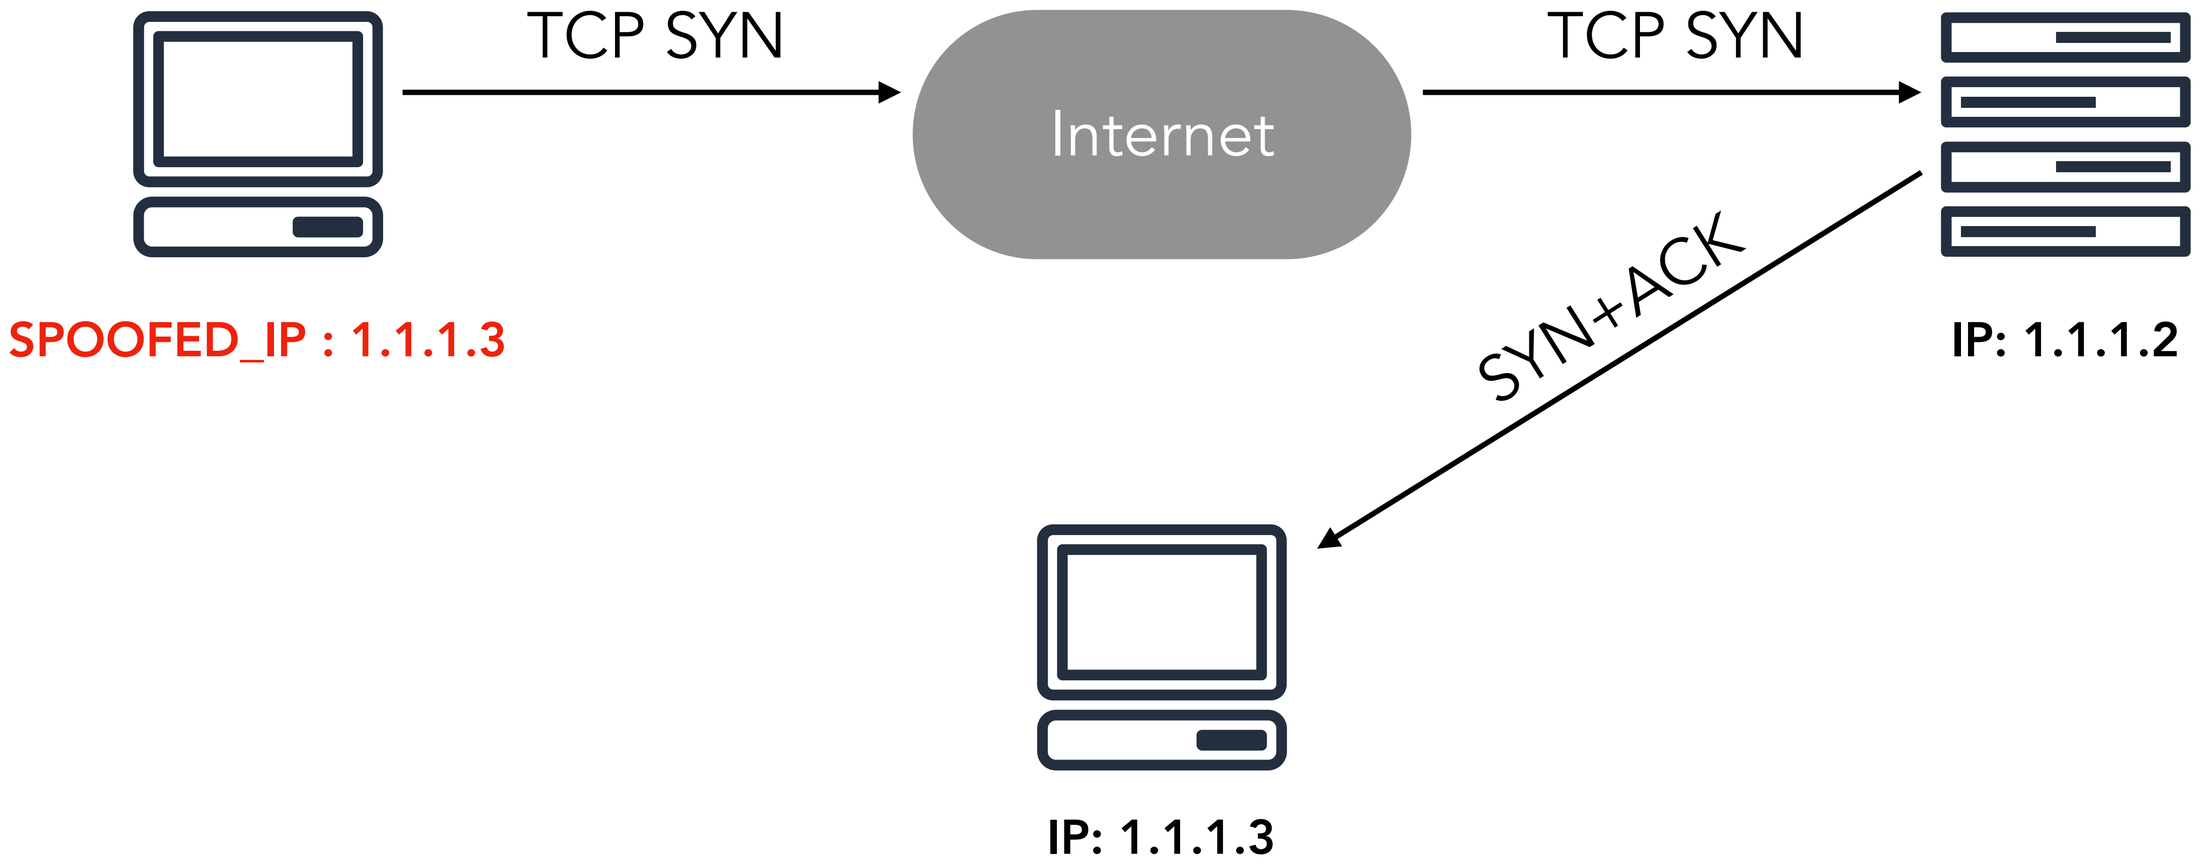

Inferring Internet Denial-of-Service Activity

A few words on this paper…

- This is a UCSD paper!
  - Stefan Savage + Geoff Voelker are the faculty authors – they do awesome work in all things cybersecurity
- This paper won best paper at USENIX Security 2001
- This paper won the USENIX Security Test-of-Time award in 2017
- First ever quantitative experiments measuring DDoS… the technique sort of started an entire field

Basic Premise of the Measurement

- Most DoS and DDoS attacks employ random IP spoofing to send attack packets. Why?

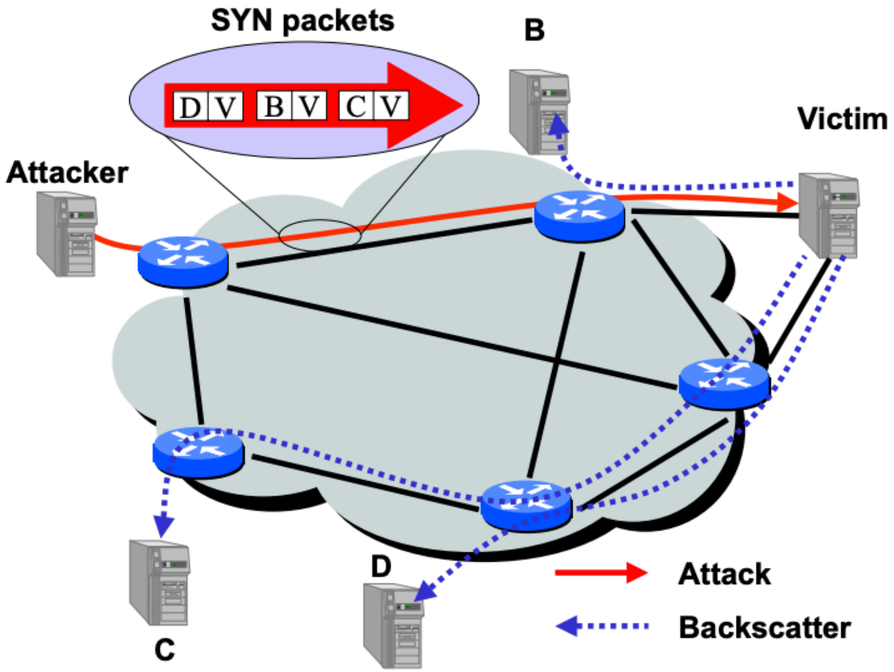

Basic Premise of the Measurement

- Most DoS and DDoS attacks employ random IP spoofing to send attack packets. Why?
- What is “backscatter?”

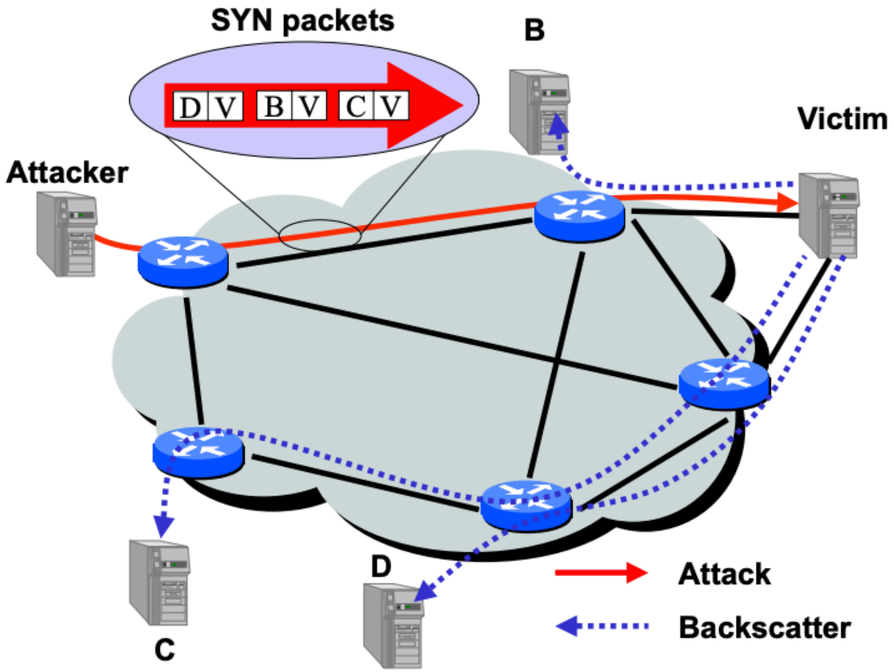

Basic Premise of the Measurement

- Most DoS and DDoS attacks employ random IP spoofing to send attack packets. Why?
- What is “backscatter?”
  - Responses sent from victim hosts and on-path devices to attack traffic

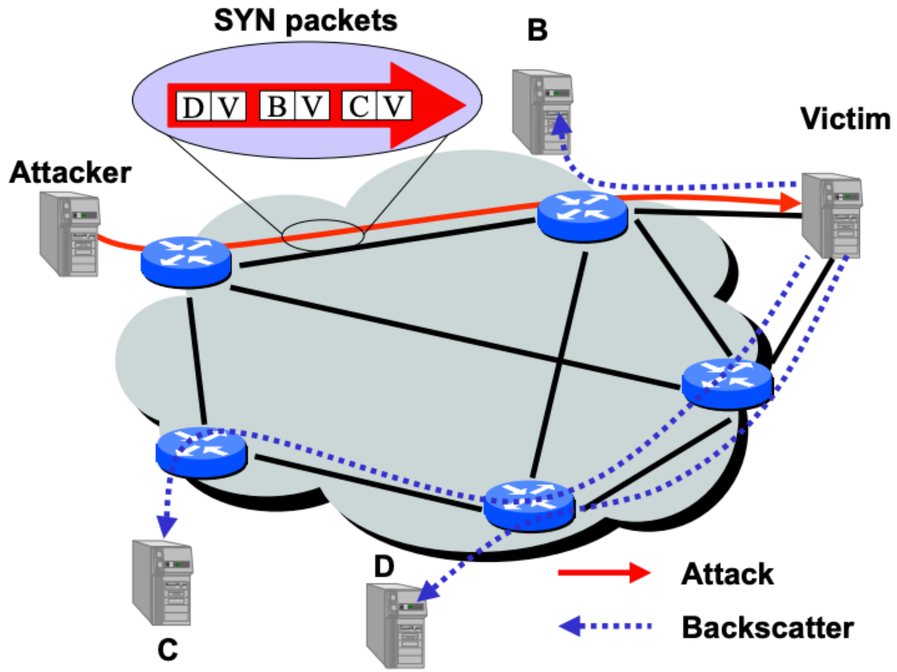

Assumptions in the paper

- What is address uniformity?

Assumptions in the paper

- What is address uniformity?
  - Why is uniformity not always guaranteed?

Assumptions in the paper

- What is address uniformity?
  - Why is uniformity not always guaranteed?
- What is reliable delivery?

Assumptions in the paper

- What is address uniformity?
  - Why is uniformity not always guaranteed?
- What is reliable delivery?
  - Why is reliable delivery not always guaranteed?

Assumptions in the paper

- What is address uniformity?
  - Why is uniformity not always guaranteed?
- What is reliable delivery?
  - Why is reliable delivery not always guaranteed?

In spite of its limitations, we believe our overall approach is sound and provides at worst a conservative estimate of current denial-of-service activity.

Experiments

- What did the authors monitor in order to identify DoS traffic?

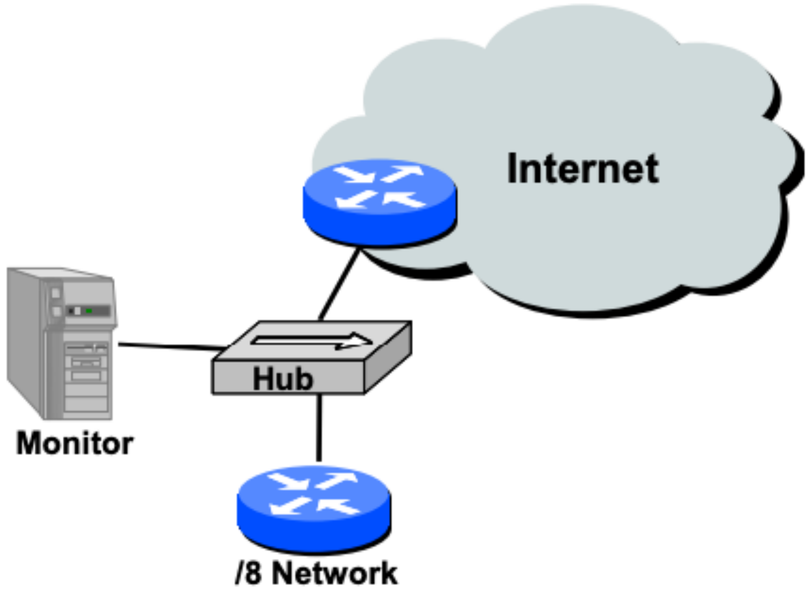

Experiments

- What did the authors monitor in order to identify DoS traffic?
  - “Slash 8” (CIDR notation), which is **1/256 of all of the Internet**
  - We call this a **darknet:** globally routable address space that goes “nowhere”
  - CAIDA operates this: https://www.caida.org/projects/network_telescope/

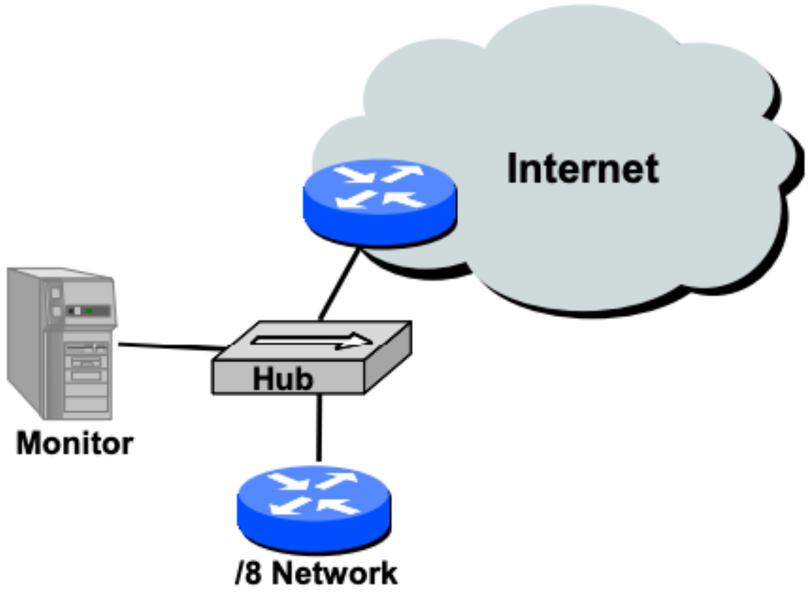

Results

- **Authors found 12,805 attacks over the course of a week,** with 5000 distinct victim IP addresses in more than 2000 distinct DNS domains
- Characterized attacks, protocols, attack durations, strategies, and domains attacked
- My highlights
  - Hackers can have a lot of fun – reverse DNS names of the victims, is.on.the.net.illegal.ly, the.feds.cant.secure.their.shellz.ca
  - Classic ***gaming*** induced DDoS: **battle.net** is a common victim
  - Tons of attacks on Internet entities (e.g., aol.com, etc.)
  - Some attacks on core infrastructure (e.g., routers that can be central points of failure)

Meta-thoughts on the paper

- DDoS remains a problem (we'll talk about Mirai shortly) – why is this still an issue? What remains the fundamental tension here that makes DDoS so hard to defeat?
- What did we like about this paper? What didn't we like?
- What could we do about DoS attacks as they are described in the paper?

Break Time + Attendance

https://tinyurl.com/cse227-attend

Codeword:
DDosferatu

Understanding the Mirai Botnet

A few words on this paper…

- This is one of my papers from 1st year of graduate school, was eventually a chapter in my PhD thesis
- Huge collaborative effort: 19 authors from 7 institutions *all* trying to understand and report on this problem
- My specific contributions in the paper:
  - Telnet honeypots
  - Malware analysis
  - Device attribution

Story time…

- I was a first year PhD student sitting in a special topics (e.g., CSE 291) security class in my first semester, when I hear Brian Krebs' website is down due to DDoS
- Krebs writes out that it's a huge attack, 620 Gbps (at the time was very large)
- Claim is that it's powered by weak IoT devices (unverified)

Story time…

- Folks in the lab think it's interesting, we want to investigate it
- Then, 9 days later, **the Mirai code is released online,** Internet havoc ensues

[FREE] World's Largest Net:Mirai Botnet, Client, Echo Loader, CNC source code release

Yesterday, 12:50 PM (This post was last modified: Yesterday 04:29 PM by **Anna-senpai**.)

Anna-senpai
L33t Member
L33T

Preface

Greetz everybody,

When I first go in DDoS industry, I wasn't planning on staying in it long. I made my money, there's lots of eyes looking at IOT now, so it
However, I know every skid and their mama, it's their wet dream to have something besides qbot.

So today, I have an amazing release for you. With Mirai, I usually pull max 380k bots from telnet alone. However, after the Kreb DDoS,
shutting down and cleaning up their act. Today, max pull is about 300k bots, and dropping.

So, I am your senpai, and I will treat you real nice, my hf-chan.

Story time…

- A few weeks later, much of the entire Internet was **down** for at least a few hours due to a Mirai DDoS attack on Dyn (dynamic DNS provider)

DDoS on Dyn Impacts Twitter, Spotify, Reddit

October 21, 2016 175 Comments

Criminals this morning massively attacked Dyn, a company that provides core Internet services for Twitter, SoundCloud, Spotify, Reddit and a host of other sites, causing outages and slowness for many of Dyn's customers.

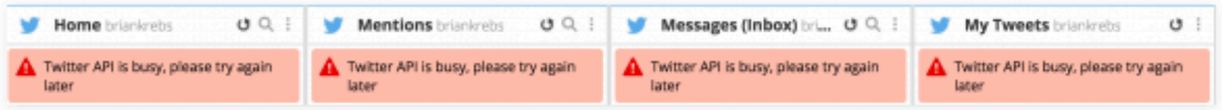

Twitter is experiencing problems, as seen through the social media platform Hootsuite.

This paper

- What is going on with the Mirai botnet?
  - How do we measure the growth, size, spread, and impact of Mirai?
- What devices enabled Mirai's power, and what was their security posture like?

How does Mirai work?

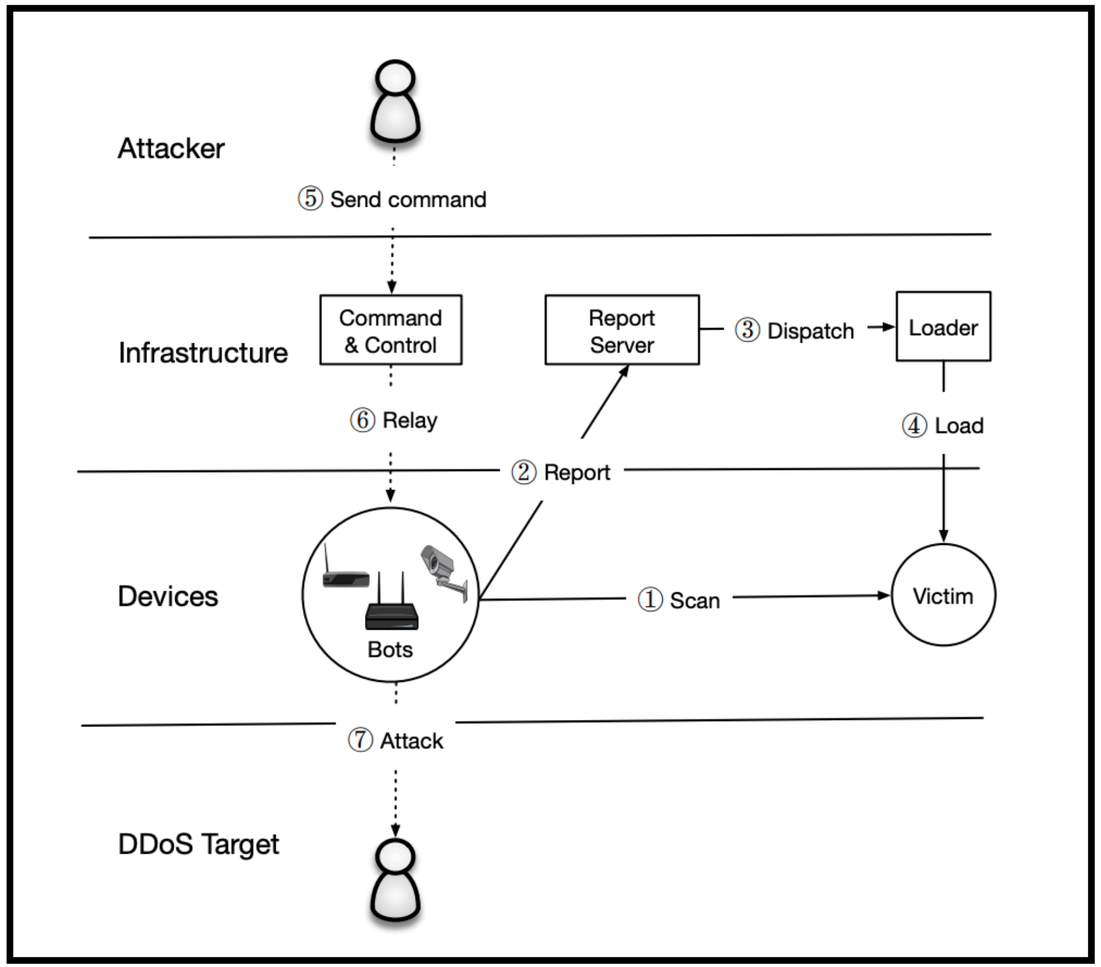

- What was the fundamental vulnerability Mirai exploited?
- What is the scanning phase?
- What is the reporting phase?
- What is the dispatch and loading phase?
- How does an attacker dispatch commands to the bots?

Methods

- We used seven different vantage points to understand the Mirai botnet
  - Network Telescope (darknet), same methods as the previous paper
  - Active scanning, same methods as the Ps and Qs paper. Why use active scanning?
  - What is a telnet honeypot? What did we use them for?
  - Passive & Active DNS
    - What is passive DNS?
  - Attack commands and attack traces from Akamai, Google Shield, Brian Krebs, Dyn, etc.

Identifying Mirai in the backscatter

- How can we distinguish Mirai scans from other DDoS traffic?

Identifying Mirai in the backscatter

- How can we distinguish Mirai scans from other DDoS traffic?
  - “Quirk” in the code that enables stateless scanning: **TCP sequence number was set to the destination IP address**
  - Expectation to see this pattern is 86 packets over our scanning period, instead we saw **116.2 billion packets**

Tracking Mirai over time – Day 0

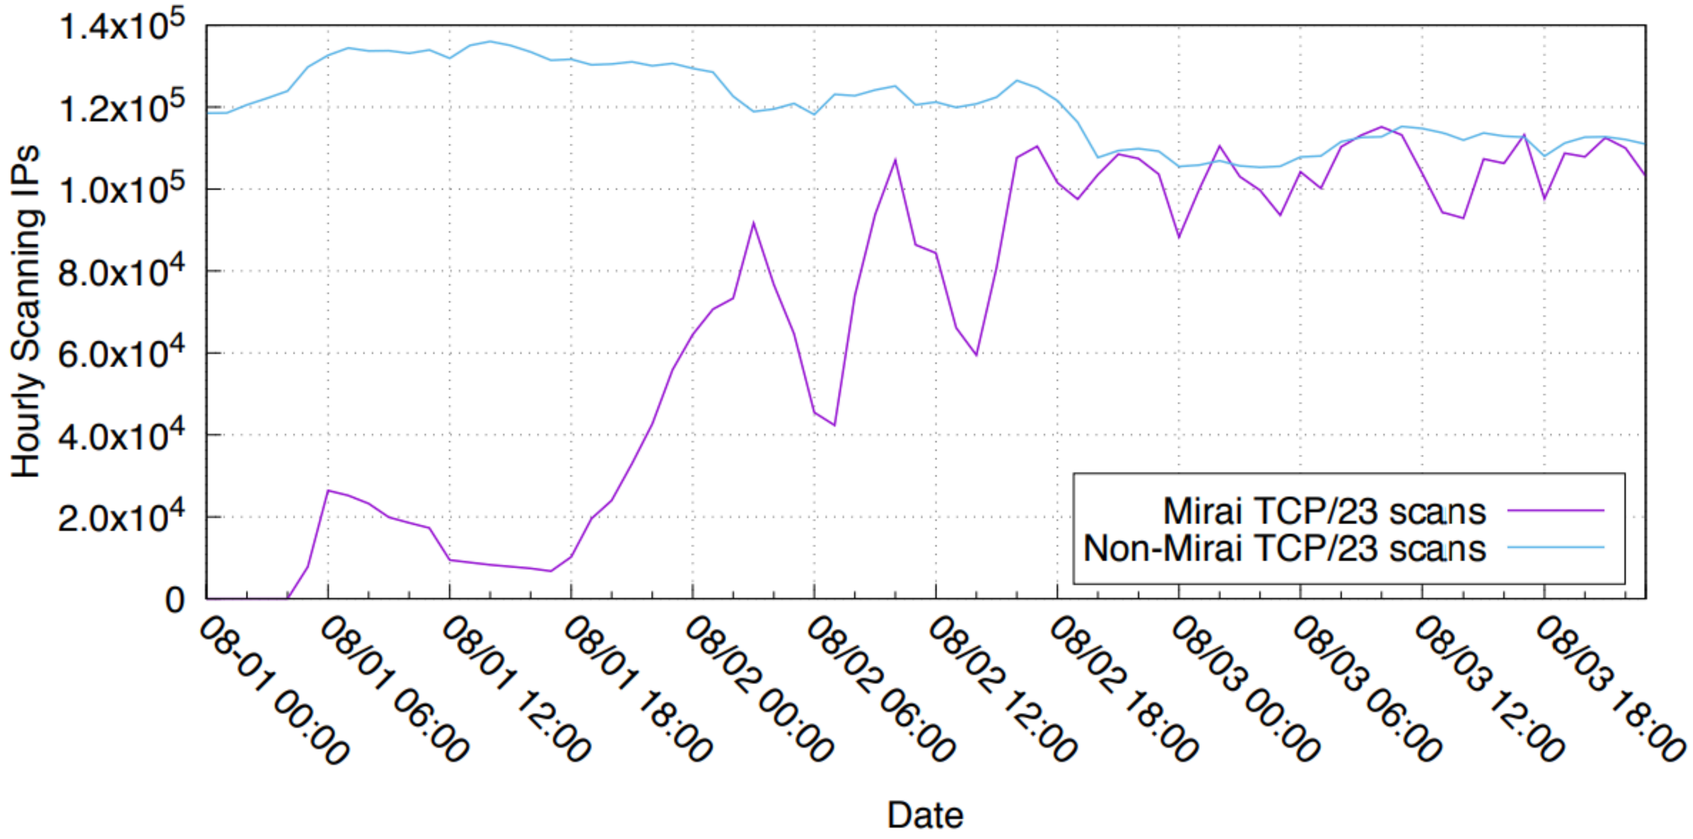

Tracking Mirai over time – Day 0

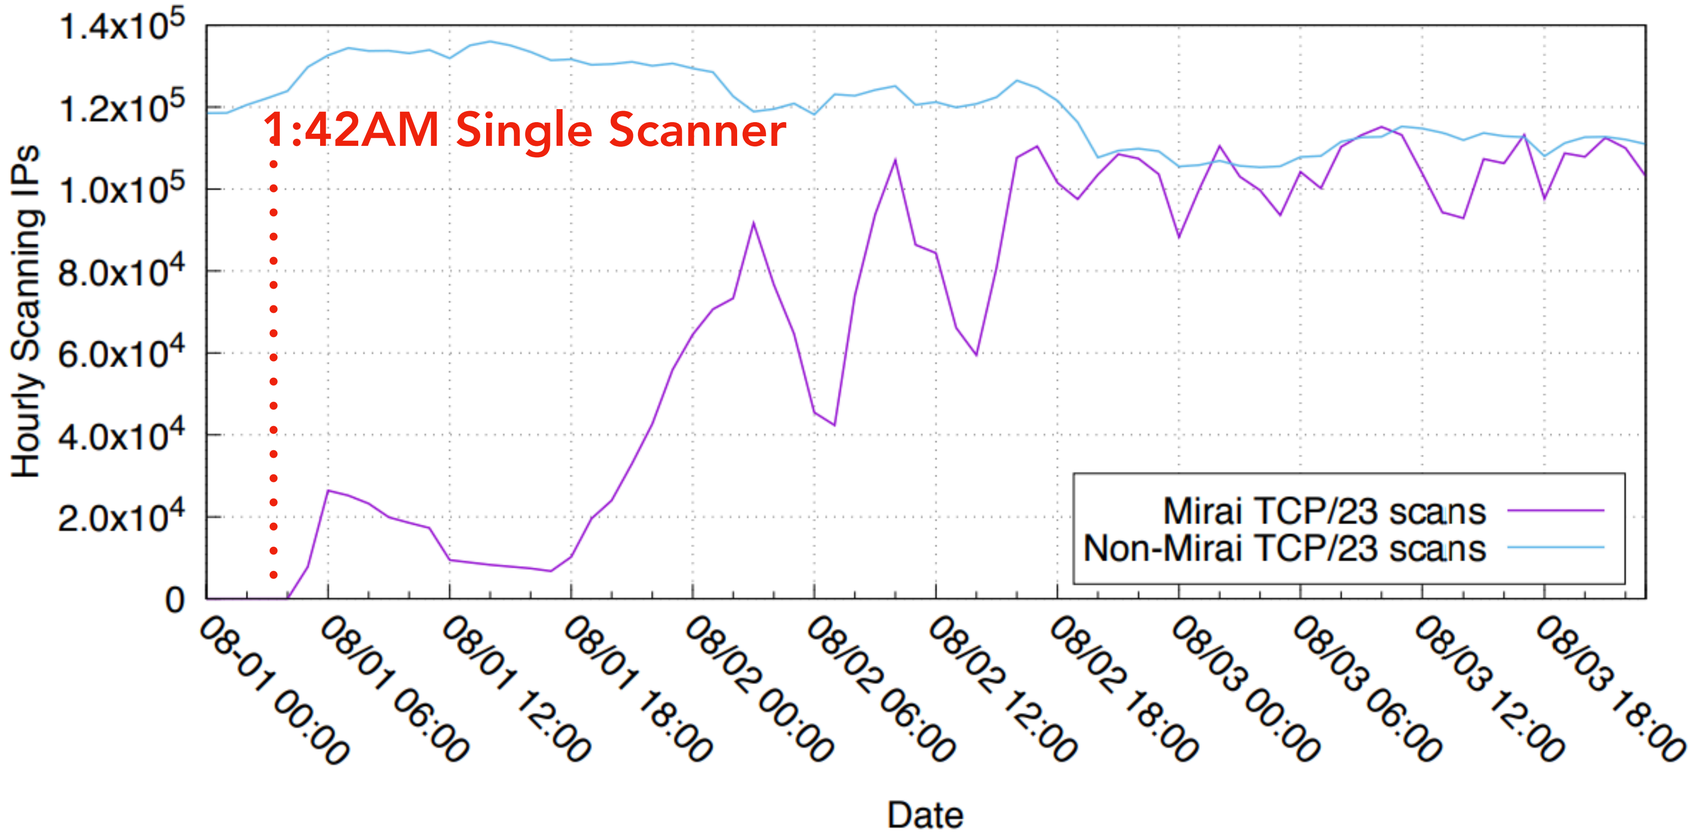

Tracking Mirai over time – Day 0

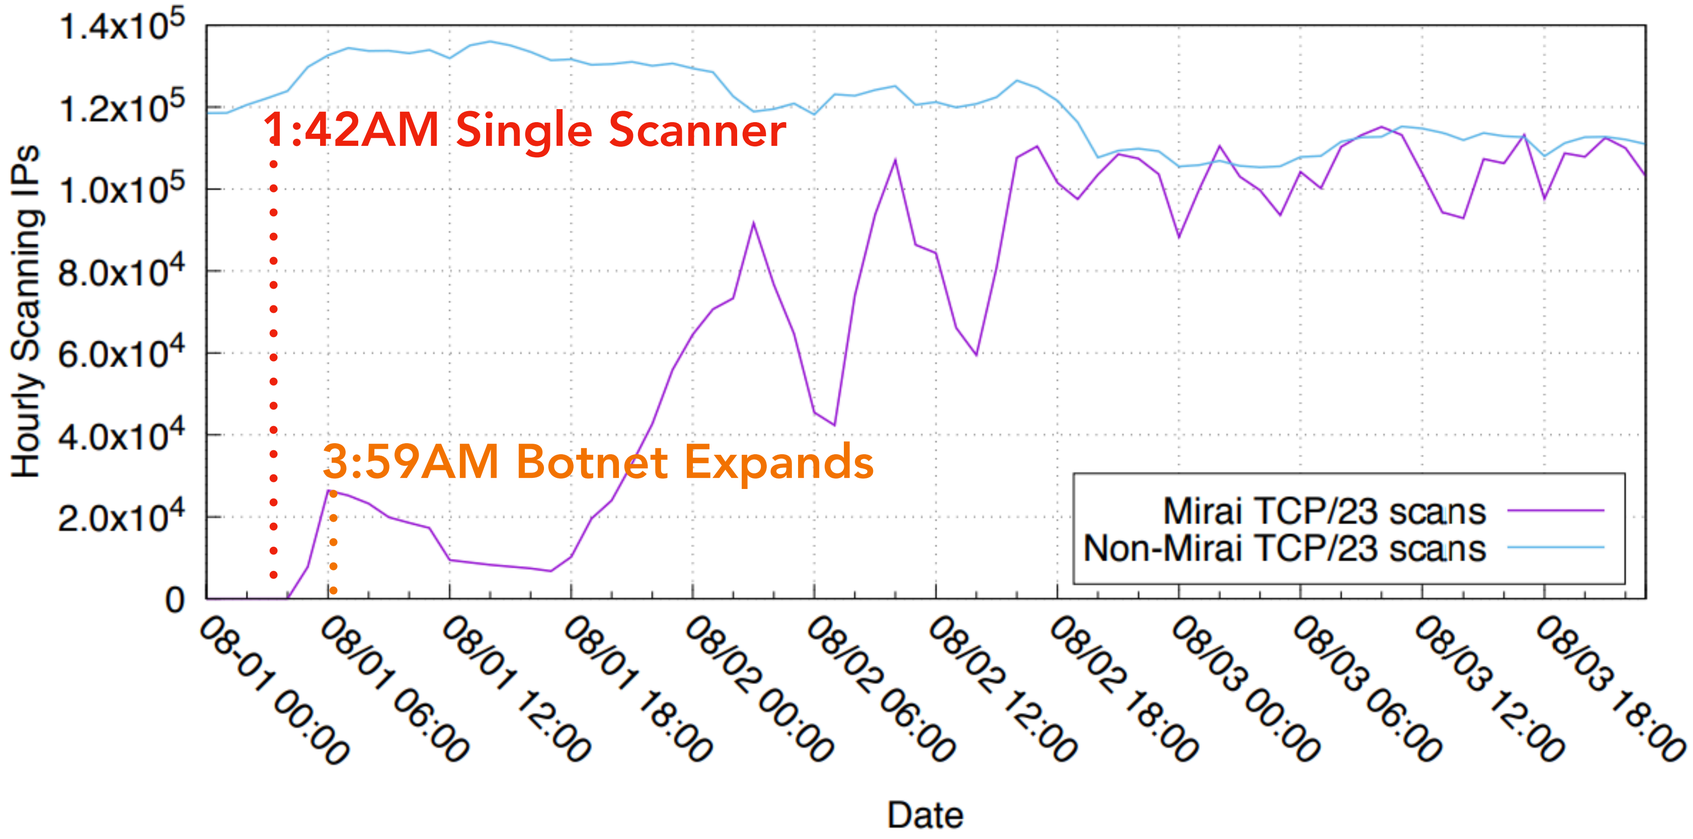

Tracking Mirai over time – Day 0

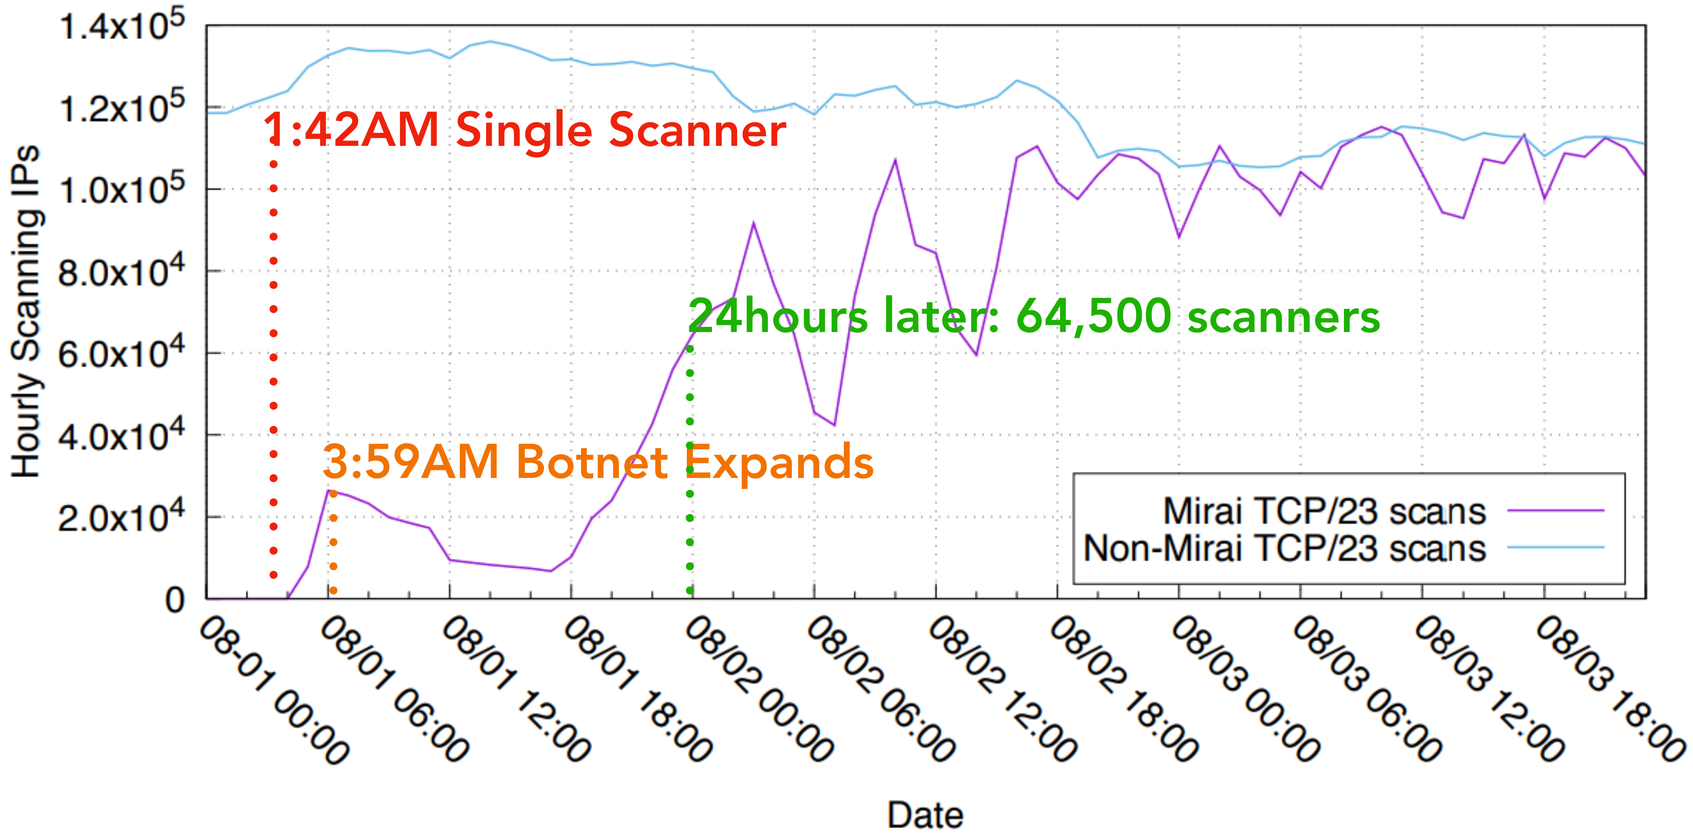

Tracking Mirai over time – the later days

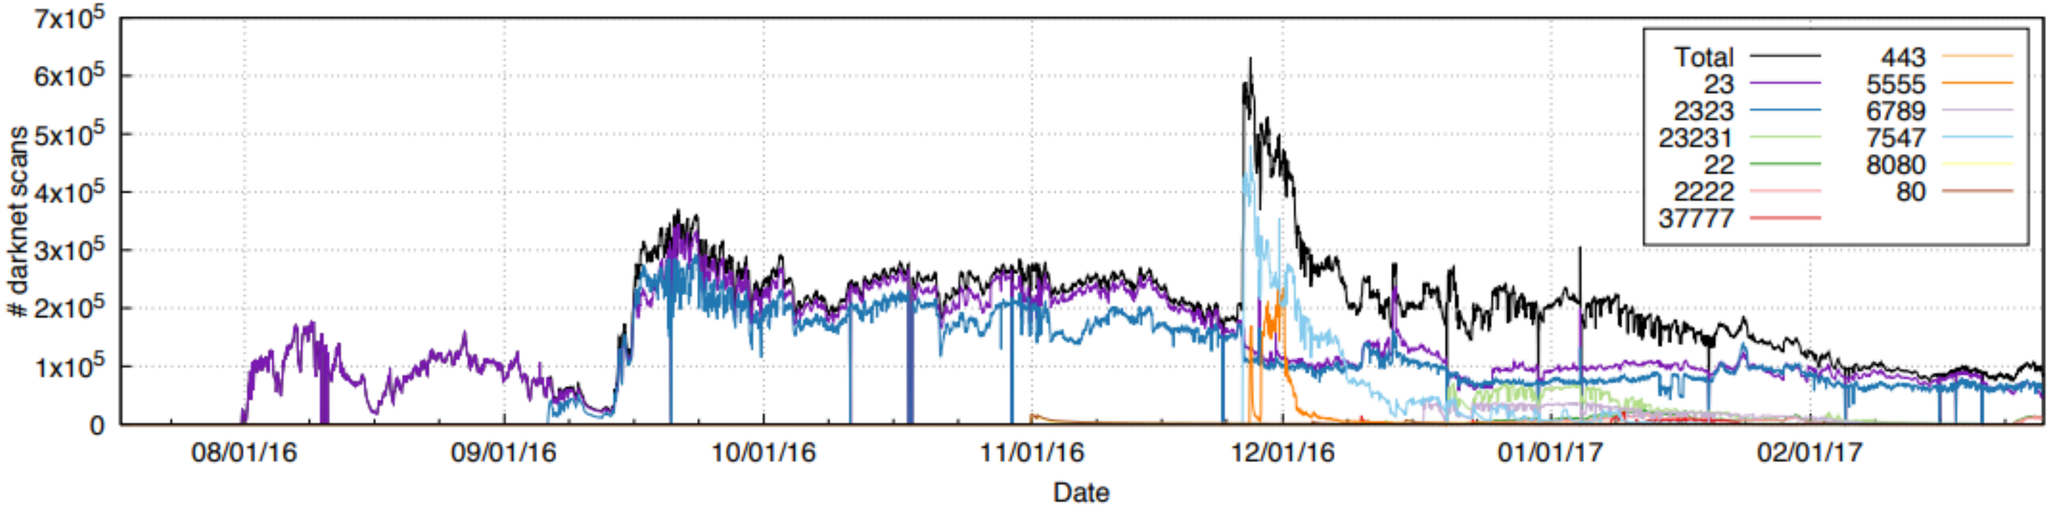

Tracking Mirai over time – the later days

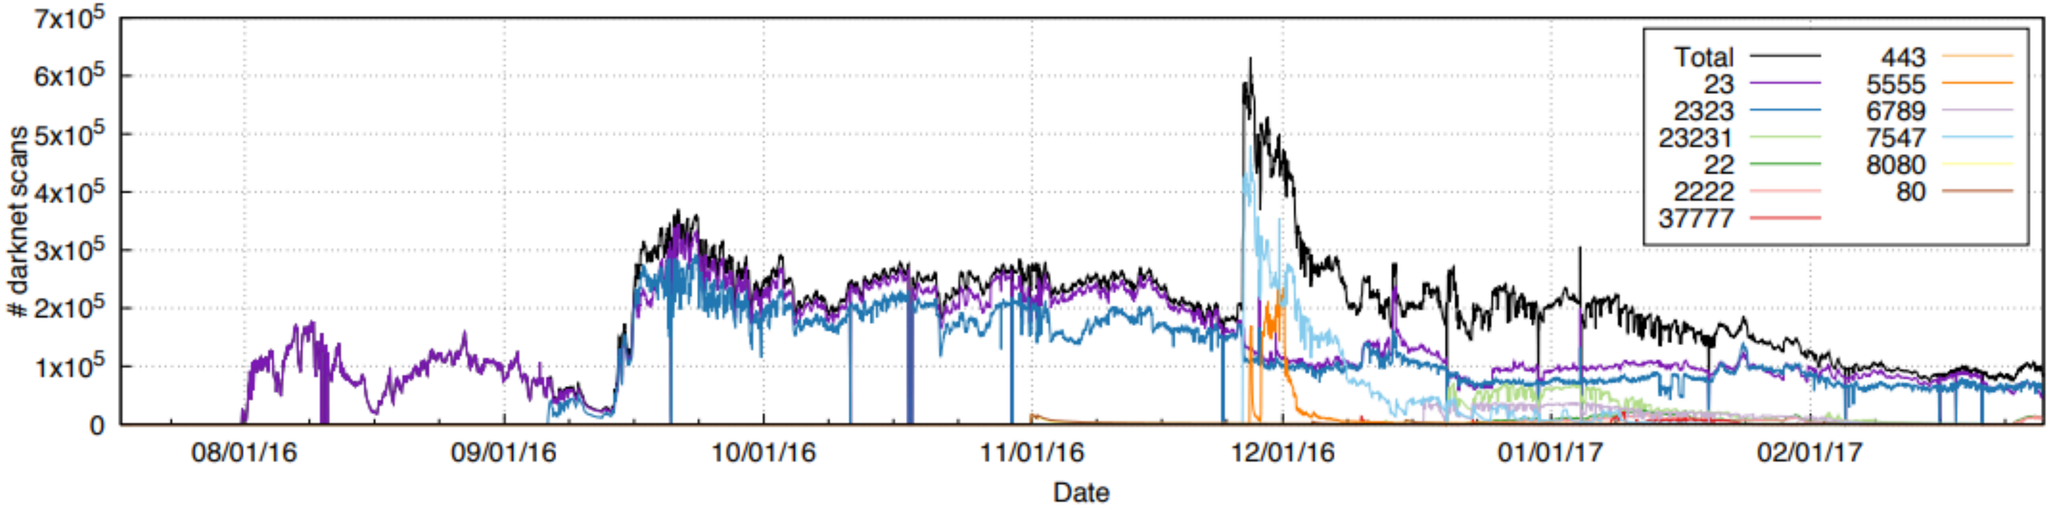

Steady state of 200 – 300K bots, with a spike in November 2016 of 600K. Why?

Where are Mirai infections happening?

Mirai

South America + Southeast Asia = 50% of Infections

TDSS/TDL4

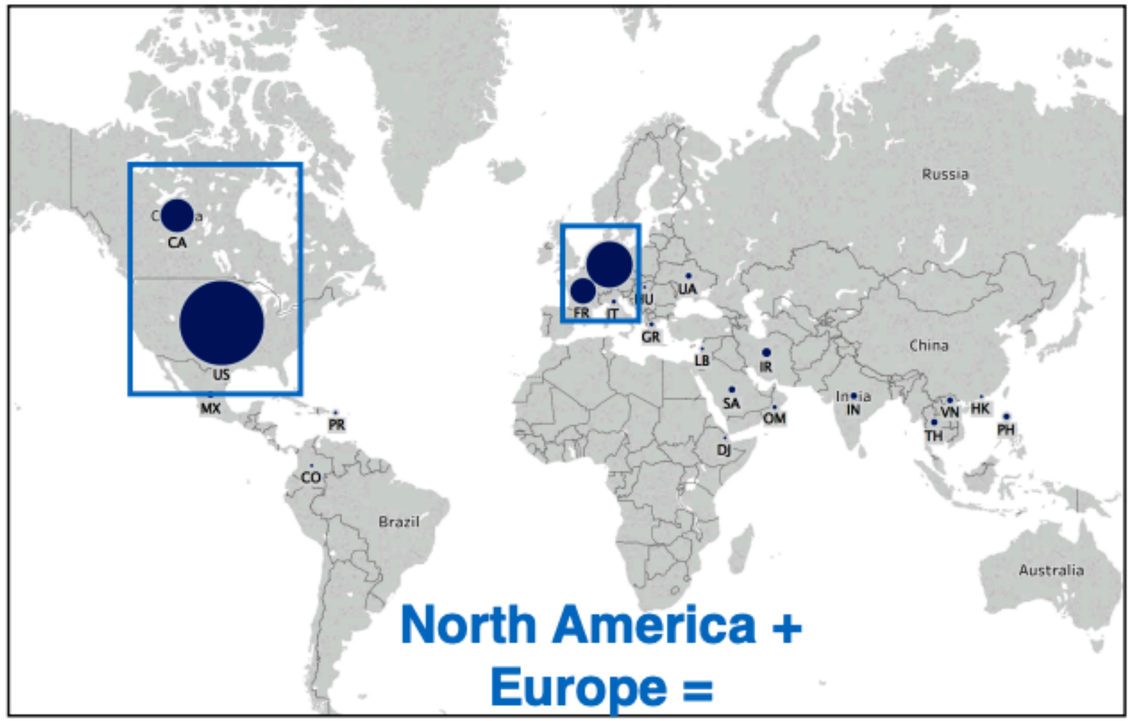

North America + Europe = 94% of Infections

How did Mirai evolve?

- How did the authors identify changes to the malware?

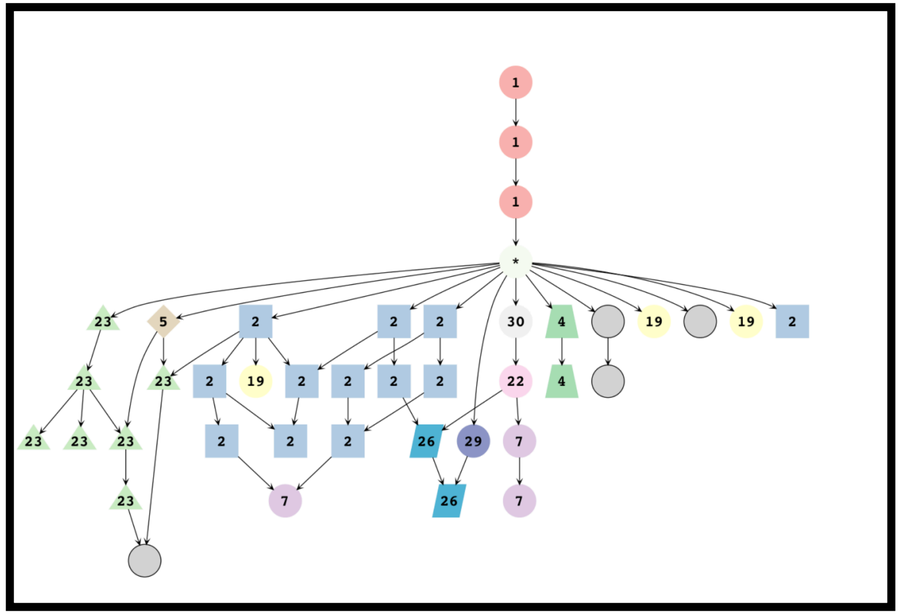

How did Mirai evolve?

- How did the authors identify changes to the malware?
- What were some of the major changes the authors identified as the malware morphed?

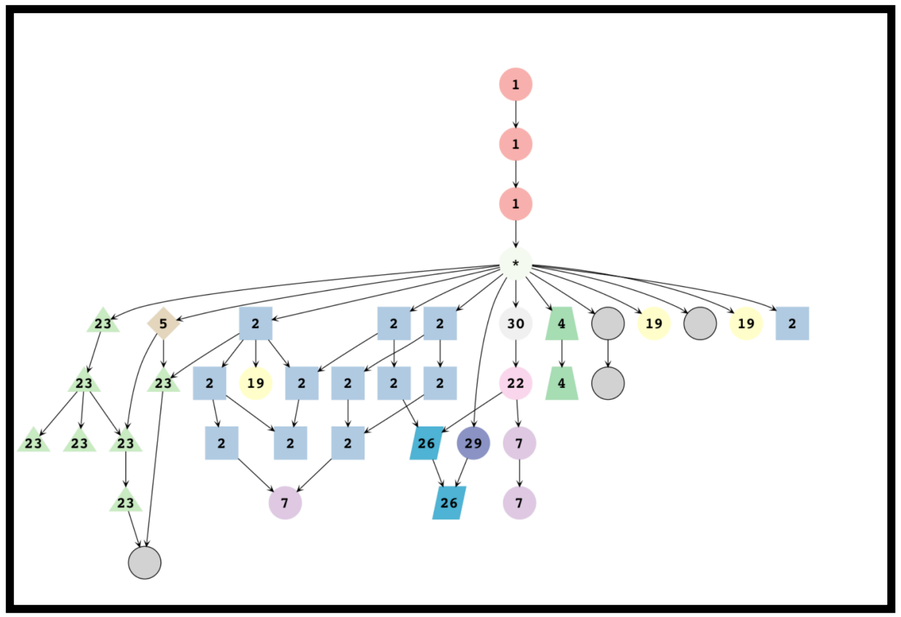

How did Mirai evolve?

- How did the authors identify changes to the malware?
- What were some of the major changes the authors identified as the malware morphed?
  - New passwords
  - New IP blocklists (e.g., removing DoD blocks from the blacklist, lol)
  - CWMP scanning

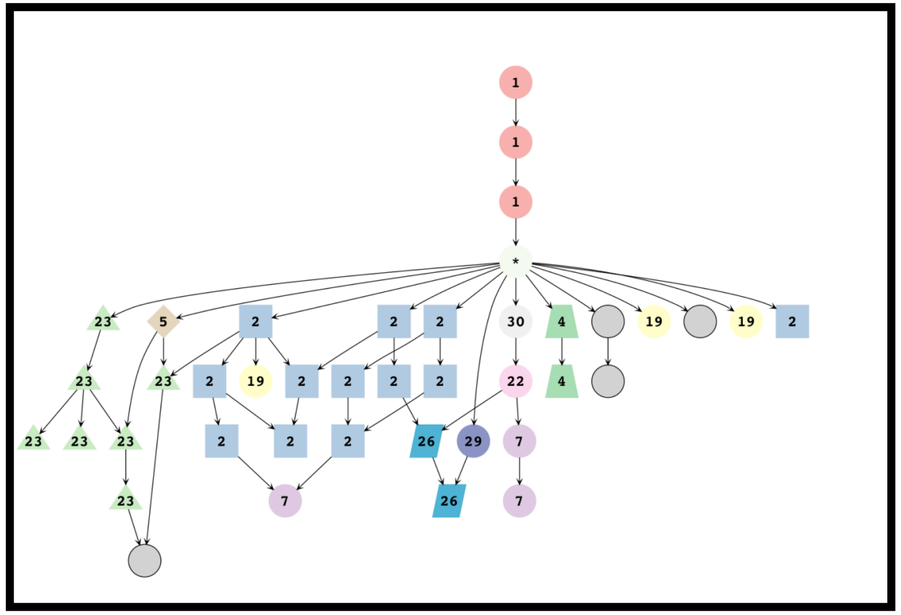

Understanding the Dyn attack

The New York Times

"It is possible, investigators say, that the attack on Dyn was conducted by a criminal group that wanted to extort the company. Or it could have been done by "hacktivists." Or a foreign power that wanted to remind the United States of its vulnerability."

Understanding the Dyn attack

The New York Times

> "It is possible, investigators say, that the attack on Dyn was conducted by a criminal group that wanted to extort the company. Or it could have been done by "hacktivists." Or a foreign power that wanted to remind the United States of its vulnerability."

Targeted IP	rDNS	Passive DNS
208.78.70.5	ns1.p05.dynect.net	ns00.playstation.net
204.13.250.5	ns2.p05.dynect.net	ns01.playstation.net
208.78.71.5	ns3.p05.dynect.net	ns02.playstation.net
204.13.251.5	ns4.p05.dynect.net	ns03.playstation.net

- Top targets are linked to Sony PlayStation
- Dyn just happened to be in the same IP block as PSN, **collateral damage**

Myriad of Targets

- **Games:** Minecraft, Runescape, etc.
- **Politics:** Chinese political dissidents, regional Italian politician
- **Anti-DDoS:** DDoS protection services
- **Misc:** Russian cooking blog….?

What happened next?

Who is Anna-Senpai, the Mirai Worm Author?

January 18, 2017 248 Comments

On September 22, 2016, this site was forced offline for nearly four days after it was hit with "**Mirai**," a malware strain that enslaves poorly secured Internet of Things (IoT) devices like wireless routers and security cameras into a botnet for use in large cyberattacks. Roughly a week after that assault, the individual(s) who launched that attack — using the name "**Anna-Senpai**" — released the source code for Mirai, spawning dozens of copycat attack armies online.

After months of digging, KrebsOnSecurity is now confident to have uncovered Anna-Senpai's real-life identity, and the identity of at least one co-conspirator who helped to write and modify the malware.

Mirai co-author Anna-Senpai leaked the source code for Mirai on Sept. 30, 2016.

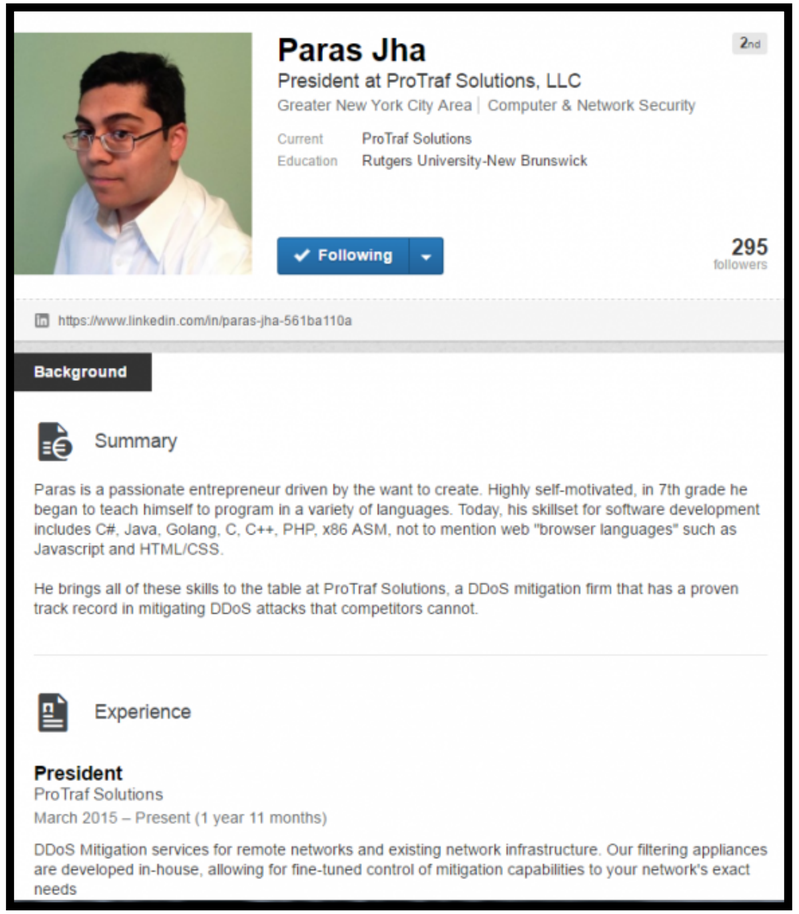

https://krebsonsecurity.com/2017/01/who-is-anna-senpai-the-mirai-worm-author/

What happened next?

The Mirai Confessions: Three Young Hackers Who Built a Web-Killing Monster Finally Tell Their Story

Netflix, Spotify, Twitter, PayPal, Slack. All down for millions of people. How a group of teen friends plunged into an underworld of cybercrime and broke the internet—then went to work for the FBI.

https://www.wired.com/story/mirai-untold-story-three-young-hackers-web-killing-monster/

Meta-thoughts on the paper

- What did we think about this paper? What did we like, what didn't we like?
- What do we do about these DDoS attacks? How do we stop them?
  - BCP 38 – http://bcp38.info/, simple idea: **ingress filtering** by the ISP, but it's not done in practice (for reasons that escape me…)
- Final thoughts?

Next time…

- All about DNS – how it works, DNS attacks, etc.
- Midpoint check-in due **Friday!**