Can ISPs Help Mitigate IoT Malware?

A Longitudinal Study of Broadband ISP Security Efforts

Arman Noroozian [1], Elsa T. Rodriguez [1], Elmer Lastdrager [2], Takahiro Kasama [3], Michel van Eeten [1], Carlos H. Gañán [1]

(1) TU-Delft (Netherlands), (2) SIDN Labs (Netherlands), (3) NICT (Japan)

The Poor State of Security in IoT Devices

- Growing influx of IoT devices into our markets
- Many new and existing devices that are poorly secured and/or raise privacy concerns
  - Poor/hard-coded passwords, lacking firmware update mechanism, using insecure communication, ...
- Many that are actively exploited (Mirai, Bashlite, Satori, Fbot, Hajime, VPNFilter, ...)
  - Used for botnets, DDoS attacks, evading detection, proxies, ...

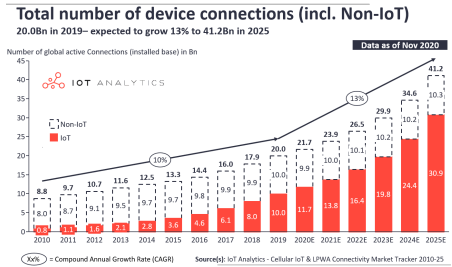

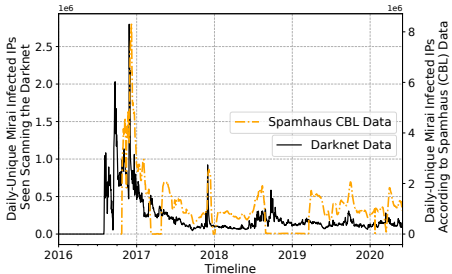

Remedies Being Explored

Many of the solutions we are exploring with respect to the IoT problem ...

- Monitoring and transparency

- Awareness raising

- Certifications and standards

- Manufacturer/vendor liability and duty to care

- Strengthening end-user rights

involve slow processes that take quite some time to take effect.
(So what can be done in the meantime?)

The Significance of Broadband ISPs

- ISPs are critical in botnet mitigation *(Asghari et al. - Post-mortem of a zombie - Usenix 2015)*
- Have security expertise
  - Fighting windows malware like *conficker* and spam botnets
- Have ability to combat infections
  - Can detect infections (e.g. Mirai)
  - Can even detect IoT devices *(Perdisci et al. - IoTFinder - IMC 2020)*
  - Are essentially gatekeepers

The Role of Broadband ISPs

- (Potentially) Have incentive to combat infections
  - 87% of (Mirai) infected IoT devices are in their networks *(Cetin et al. - Cleaning Up the Internet of Evil Things - NDSS 2019)*
  - Large number of exploited IoT devices are their own routers
- Have security practices at their disposal that are known to work against IoT malware
  - Quarantine infected networks *(Cetin et al. - Cleaning Up the Internet of Evil Things - NDSS 2019)*

ISP Security Practices to Combat IoT Malware

- Quarantining works but is costly and difficult to scale
- Can ISPs effectively help with mitigating IoT malware through security best practices beyond quarantining?

- We examine two additional strategies
  - **Does reducing the IoT attack surface help?**
  - **Do existing abuse remediation practices also work against IoT Malware?**

Study Setup (Scope)

- We limit study to *Mirai* and *Mirai-like* IoT Malware because they are suitable as case-studies for several reasons:
  - Mirai is (still) among major IoT malware families, numerous IoT malware families are based on its source code
  - Easier to track
    - Aggressive scanning to infect more devices
    - Unique network traffic fingerprint (scan packets with $TCP_{seq} = DEST_{IP}$ *Antonakakis et al. - Understanding Mirai - USENIX 2017*)
  - Easier to cleanup
    - power-cycling device removes Mirai (non-persistent)
  - Abuse / Threat Intel data on Mirai largely available and shared with ISPs (Shadowserver, Spamhaus ...)
  - If ISPs cannot mitigate Mirai and Mirai-like malware, unlikely they will be able to do so for more sophisticated IoT malware

Study Setup (Method)

- Model how changes in number of infections within each broadband operators' network(s) correlate with:

1. IoT attack surface reduction in the network
   (a) Measured by proxy of changes to number of accessible ports that Mirai uses for propagation: eg. `TCP/21,23,2323,7547`
2. Network hygiene and abuse remediation efforts of the operators
   (a) Measured by proxy of changes to hygiene indicators (number of open DNS resolvers that can be exploited for DRDoS attacks, number of non-IoT, and other non-Mirai IoT infections)

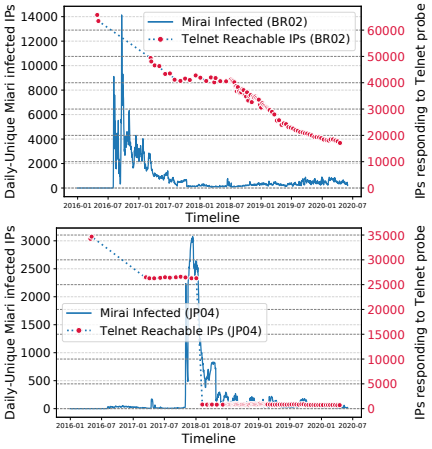

Study Setup (Data)

- Longitudinal study of *Mirai-(like)* infections in global Broadband ISP networks
  - Collect darknet data from January 2016 to May 2020 (over 4 years) *(Antonakakis et al. - Understanding Mirai - USENIX 2017)*

Overview of data collected and used in our study

Period(s) Covered	Description	Source	Type
2015, 2019	Broadband ISP network mapping and statistical data	`www.telegeography.com`	Commercial/Marketing
2016-04 ↔ 2020-04	Probes of various TCP services (TCP/21,23,53,2323,7547)	Rapid7 Project Sonar (TCP Scans + National Exposure Scans)	ZMAP scans of IPv4 space
2016-01 ↔ 2020-05	IPs with Mirai-like infections	NICT (Japan)	Darknet data
2016-10 ↔ 2020-05	IPs with non-IoT botnet infections	Spamhaus (CBL)	Anti-Abuse / Threat-Intel Feed
2016-10 ↔ 2020-05	IPs with non-Mirai IoT malware infections	Spamhaus (CBL)	Anti-Abuse / Threat-Intel Feed

Some Empirical Results

- Mirai-like infections moderately to strongly correlate with the number of ISP subscribers
- Also find moderate correlations with other factors
  - Attack surface: Number of reachable FTP, Telnet, TR069 services
  - Network Hygiene: Number of *other* malware infections
- We find evidence that many ISPs and their users reduced the attack surface and as well as variations in broadband network hygiene over time

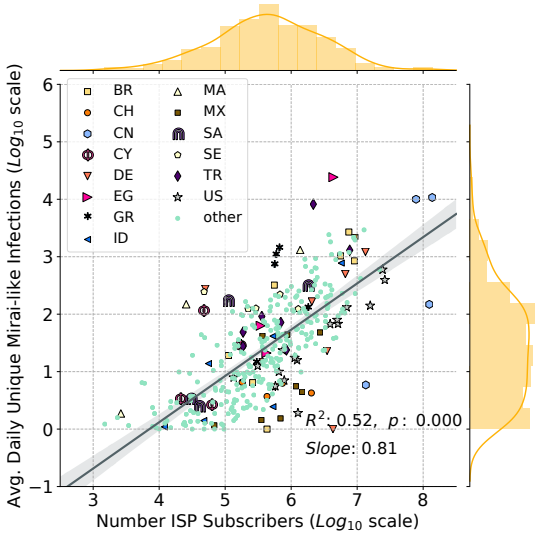

Regression Modeling

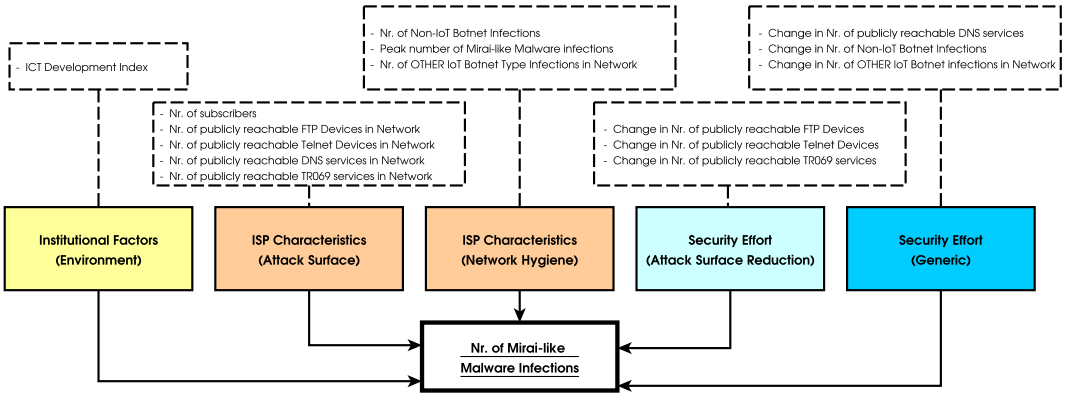

Modeling Results

- Does attack surface reduction correlate with less Mirai-like infections after controlling for other factors?
  - Surprisingly we found no evidence to empirically support that.
- Does improved network hygiene correlate with less Mirai-like infections then?
  - Yes, we find that broadband networks that have poorer network hygiene and abuse mitigation also have higher infection rates for Mirai and vice versa.

Takeaways

1. A lack of evidence to support attack surface reduction being effective may be explained by newer Mirai variants having moved on to alternative exploitation and propagation methods which we didn't account for (A whack-a-mole phenomenon)
   (a) Does not suggest that attack surface reduction is a lost endeavor
   (b) Would have surely seen higher infection levels without it
2. Overall we do find evidence to support that ISPs may play a significant role in combating IoT malware.
   (a) Best practices for general botnet mitigation appear to also be relevant for IoT malware.

Takeaways (cont.)

3. ISPs have several countermeasures at their disposal
   (a) Better protecting customer networks
      i. e.g. via more secure default configurations on router equipment
      ii. Ports that are closed by default
      iii. Stronger initial passwords
      iv. Firewall rules that prevent mass scale port scanning
   (b) Abuse Handling
      i. Notifying infected customers
      ii. Quarantining infected customer networks
      iii. Updating equipment and their firmware

4. The role of ISPs in mitigation should not obscure the need to develop policies that tackle root cause of problem: **poor security practices of IoT manufacturers**.

Contact

✉ a.noroozian _at_ {tudelft,uva}.nl

@anoroozian

https://anoroozian.nl

Can ISPs Help Mitigate IoT Malware? A Longitudinal Study of Broadband ISP Security Efforts

Arman Noroozian[1], Elsa Turcios Rodriguez[1], Elmer Lastdrager[1], Takahiro Kasama[2], Michel van Eeten[1], and Carlos H. Gañán[1]
[1]TU-Delft (Netherlands), [2]SIDN Labs (Netherlands), [3]NICT (Japan)

Abstract—**For the mitigation of compromised Internet of Things (IoT) devices we rely on Internet Service Providers (ISPs) and their users. Given that devices are in the hands of their subscribers, what can ISPs realistically do? This study examines the effects of ISP countermeasures on infections caused by variants of the notorious Mirai family of IoT malware, still among the dominant families. We collect and analyze more than 4 years of longitudinal darknet data tracking Mirai-like infections in conjunction with threat intelligence data on various other IoT and non-IoT botnets across the globe from January 2016 to May 2020. We measure the effect of two ISP countermeasures on Mirai variant infection numbers: (i) reducing the attack surface (i.e., closing ports that are used by the malware for propagation) and (ii) ISPs increasing their general network hygiene and malware removal efforts (as observed by proxy of the remediation of infections of other families of IoT and non-IoT malware and reductions in the number of DDoS amplifiers in their networks). We map our infection data to 342 broadband providers that have the bulk of the broadband market share in their respective 83 countries. We find that the number of infections correlates strongly with the number of ISP subscribers (R^2=0.55). Yet, infection numbers can still vary by three orders of magnitude even for ISPs with comparable subscriber numbers. We observe that many ISPs, together with their subscribers, have reduced their attack surface for IoT compromise by blocking traffic to commonly-exploited infection vectors such as Telnet and FTP. We statistically estimate the impact of these reductions on infection levels and, counter-intuitively, find no significant impact. In contrast, we do find a significant impact for improving general network hygiene and best malware mitigation practices. ISPs that were more successful in reducing DDoS amplifiers and non-Mirai malware infections in their networks also end up with significantly lower Mirai infection rates. In other words, rather than investing in IoT-specific countermeasures like reducing the attack surface, our findings suggest that ISPs might be better off investing in general security efforts to improve network hygiene and clean up abuse.**

Index Terms—**Mirai, Internet of Things, IoT, Malware, ISP, Countermeasure, Remediation**

During this period, millions of devices were compromised by the Mirai malware family [1]. Mirai not only caused the first peak of infections, but it has persisted as a dominant malware family until today. One recent industry report named Mirai the "king of IoT malware" [2]. In 2019, Kaspersky reported that Mirai is still the leading malware family and responsible for 21% of the IoT infected devices [3]. What keeps Mirai a relevant threat is that it exploits default credentials, a problem that has still not been fixed by many manufacturers. The Open Web Application Security Project (OWASP) describes this as the top threat for IoT [4]. Additionally, the release of Mirai's source code has allowed attackers to add exploit code on top of its credential-based attacks and create newer variants which go beyond launching Denial of Service (DoS) attacks. According to industry reports, compromised IoT devices have been abused for purposes ranging from DoS attacks to the installation of tor nodes, packet sniffers, and trojans, all the way to performing crypto-jacking, DNS hijacking, and credential collection [5].

As long as manufacturers keep releasing new insecure devices into our markets, the brunt of remediating infected IoT lies with both the end users who own the devices and the Internet Service Providers (ISPs) where more than 80% of the devices are located [6]. While recent work has studied the practices and perceptions of end users when it comes to IoT security [7–9], little attention has been paid to the role of ISPs. One exception is a study that found ISPs to be able to use quarantining of infected devices as a way to enforce remediation by the customer [6]. We know virtually nothing however, of what other ISP practices might be effective.

In this paper, we explore two additional security strategies. First, can ISPs stem the spread of IoT infections by reducing the attack surface for the malware? In other words, does it help to close network ports that are used for propagation? ISPs can administer default router configurations that block ports of commonly exploited services such as Telnet and FTP [10]. This is similar to past approaches in mitigating spam, where port 25 would be blocked to prevent the distribution of spam from consumer connections. Second, are general ISP security measures for network hygiene and abuse remediation also effective