

Malware and IPv6 Scanning

Vivian Band, University of Glasgow

Email: v.band.1@research.gla.ac.uk

Twitter: @ArcStatic42

Finding a PhD Project is Hard.

What Will Malware Look Like in the Future?

IoT devices are becoming more and more popular.

Botnets, like Mirai, often scan the IPv4 address space for mass recruitment of IoT devices which are hard to infect through social engineering.

What if Everyone Migrated to IPv6?

What if we achieve the dream and finally have a fully IPv6-enabled Internet?

How would self-propagating malware adapt when exhaustive address space scans are no longer feasible?

Reducing the IPv6 Address Space

Reducing the IPv6 Address Search Space

IPv6 addresses are split into two halves: network prefix and interface identifier (IID).



Local Network Scans

A machine on a given local network would already know the full assigned network prefix.

Network Prefix (64 bits)



Local Network Scans

A machine on a given local network would already know the full assigned network prefix.

Bytes 3 and 4 of IID are always 0xfffe.




Local Network Scans

A machine on a given local network would already know the full assigned network prefix.

Bytes 3 and 4 of IID are always 0xfffe.

Organisationally Unique Identifiers (OUIs) are associated with a manufacturer - limited list of valid OUIs.




Local Network Scans

A machine on a given local network would already know the full assigned network prefix.

Bytes 3 and 4 of IID are always 0xfffe.

Organisationally Unique Identifiers (OUIs) are associated with a manufacturer - limited list of valid OUIs.

Address space reduced from 128 to 24 bits for a local scan.



Scale of a Local IPv6 Scan

18,446,744,073,709,551,616 possible IID combinations (2^{64}).

16,777,215n possible IID combinations with reduced address space ($(2^{24})n$).

4,294,967,296 possible IPv4 addresses (2^{32}) can be scanned in around 5 minutes - 256 OUIs can be searched for in an equivalent amount of time.

Privacy Addresses

Privacy addresses are intended to obscure host addresses from passive observers.


However, according to RFC 7707, these seem to be an alias for the stable SLAAC address. Privacy addresses do not do much to mitigate active scanning attacks.

This implies that IPv6 malware samples will generate network traffic scanning for particular interface identifiers - seeing which type(s) of device and ranges a sample is targeting can provide additional clues about the aims of new malware.

Internet-Wide Scans

Active network prefixes can be obtained through Border Gateway Protocol (BGP) advertisements.

/32 advertisement gives a prefix with the first 32 bits set, /48 the first 48 bits set, etc.




Internet-Wide Scans

Active network prefixes can be obtained through Border Gateway Protocol (BGP) advertisements.

/32 advertisement gives a prefix with the first 32 bits set, /48 the first 48 bits set, etc.

/48 is the most common - let's take that as an example.



Internet-Wide Scans


Active network prefixes can be obtained through Border Gateway Protocol (BGP) advertisements.

/32 advertisement gives a prefix with the first 32 bits set, /48 the first 48 bits set, etc.

/48 is the most common - let's take that as an example.

2^{16} possible network prefixes per advertisement.

There are $(2^{24})n$ possible addresses on the network, once a valid prefix is found.



Scale of an Internet-Wide Scan

Network prefix scans can be considered separate from 24-bit host scans described earlier - $(2^{24})n$ attempts at host finding are unnecessary if a network prefix is not valid or otherwise unresponsive.



/48 IPv6 prefix advertisements are the most common, followed by /32. These leave 16-bit and 32-bit spaces to scan, respectively.





Domestic ISPs often advertise /56 prefixes. This leaves a network search space of 8 bits, which is trivial.


Malware Infection Vectors


Weakest Link in the Chain

Single devices can be infected through social engineering (phishing, fake downloads, scareware, etc) - maybe several devices, if multiple sessions are active on a compromised login-based service.

Yahoo Mail Lottery. Spam x  

 **ANNIVERSARY AWARD <wembamrrichard@gmail.com>** 11 Jun 2020, 15:06   
to bcc: me ▾



This message seems dangerous 

 Many people marked similar messages as phishing scams, so this might contain unsafe content. Avoid clicking links, downloading attachments or replying with personal information.

PAYMENT MANAGER APPROVED CONTACT INFORMATION:
FAX: +27 86 666 0488
Mobile : +27 83 673 0036
CONTACT PERSON: Mr. Trevor D. Williams
E-MAIL: williams_D_trevor@post.com
E-MAIL: mr.williams.d.trevor@gmail.com

You are advised to send the following information to your Claims Agent to facilitate the release of of your fund to you.

Full Name.....
Country.....
Contact Address.....
Telephone Number.....
Fax Number.....
My Date of birth.....
Occupation.....
Please note that a copy of your passport or your drivers license is needed while sending your information.....

 **Downloading this attachment is disabled.** This email has been identified as phishing. If you want to download it and you trust this message, click 'Not spam' in the banner above. 

Weakest Link in the Chain

Malware on more complex devices initially spreads through social engineering, which is then augmented through local network scanning.

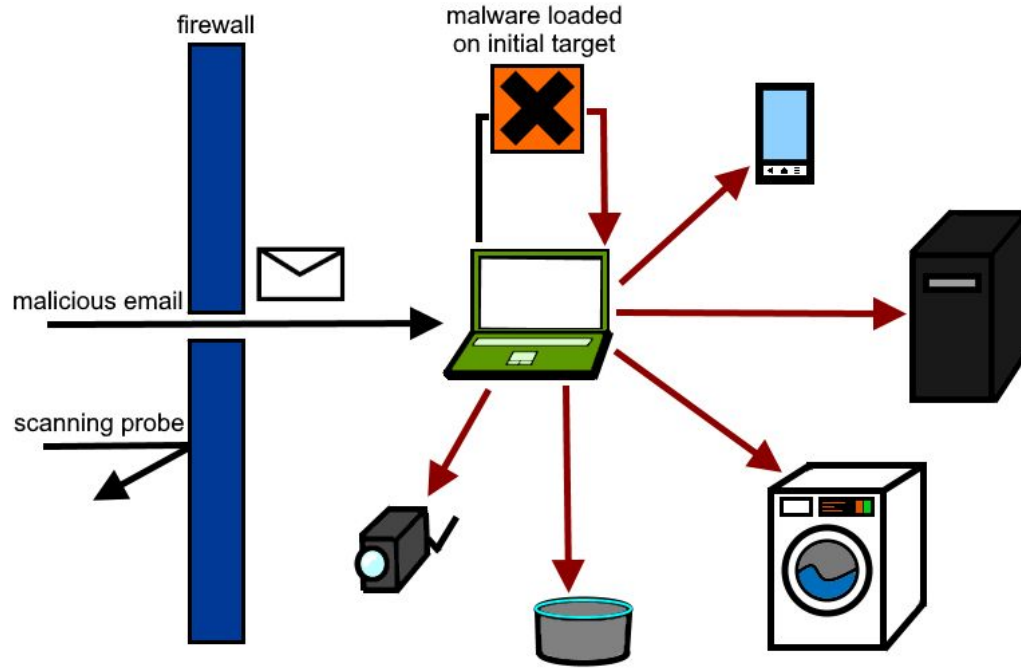
Emotet, a trojan often used to load other malware, uses brute force attacks to infect other systems - but it needs to locate the hosts first.

Easy with IPv4, difficult with IPv6.



Image: https://en.wikipedia.org/wiki/Trojan_Horse

Malicious Local Scan




Malicious Local Scan

Taking Emotet as an example: IPv6 network prefix is obtained from initial infection through a suspect attachment.

The Emotet operator wants to scan for n different OUIs known to be used in devices with a particular vulnerability.

There are $(2^{24})n$ possible addresses to scan.



Malicious Internet-Wide Scan

IoT botnets, like Mirai, very commonly use Internet-wide scans to recruit large numbers of vulnerable, Internet-facing devices.

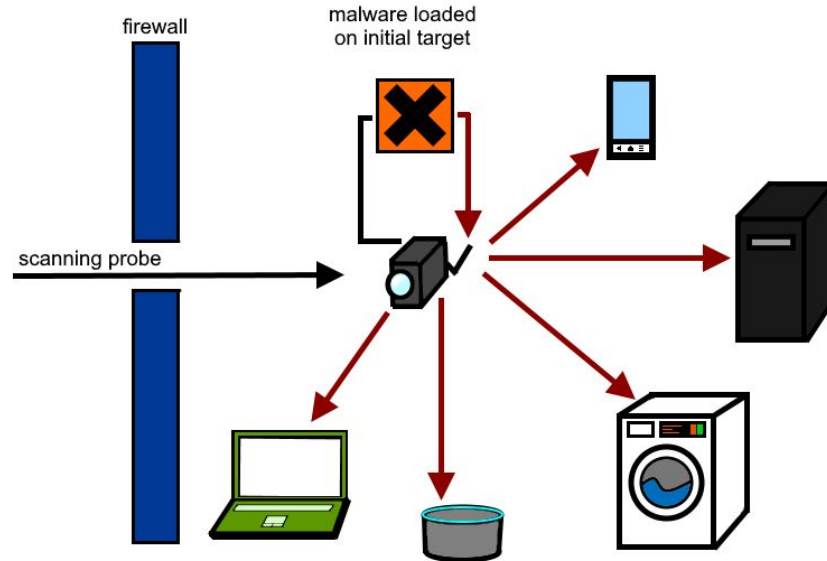
Exhaustive scan of all possible 32-bit IPv4 addresses is feasible, not possible with 128-bit IPv6 addresses.



Image:

<https://www.smithsonianmag.com/smart-news/lighted-escape-hatches-could-help-little-fish-flee-trawlers-nets-23876895/>

Malicious Internet-Wide Scan




Malicious Internet-Wide Scan

Local scans can be reduced to 24 bits, if MAC address-based IPv6 addresses are assigned.

RFC 7707 claims support for allocating these address types is mandatory.


Internet-wide scans for specific devices are split into two halves: between 32 and 8 bits for the network prefix, then 24 bits to search for specific hosts once a responsive network prefix is found.



Malicious Internet-Wide Scan

All the information used to minimise the searchable space to these ranges is publicly accessible.

IPv6 scans are entirely feasible for an attacker to perform in a feasible amount of time.



Malware Business Models

Malware-as-a-Service (MaaS)

MaaS operators lease devices to other malware operators for a fee.

Clients gain access to recruited devices through loader malware installed by MaaS organisations.

An efficient IPv6 scanning algorithm would help increase profits through recruiting more devices, and being able to offer more bespoke services.



Weaponised Data Disclosure

Some groups, like Maze, are moving towards pressuring victims into paying the ransom by threatening to put private data in the public domain.

Mostly business databases, but could evolve to be used against individuals.

Lateral movement and network recon often used here - IPv6 scans useful for both of these.



Open Source DDoS

Compared to ransomware, IoT botnet variants like Mirai are not hugely profitable.

Mostly used to DDoS high-profile targets and assemble proxy servers.

These types of malware specifically exclude certain IPv4 address ranges from search - IPv6 scans could allow them to be even more fine-grained with this.



SALE!

Botnet Package 4

~~\$80.00~~ **\$100.00**

- 6x VPS
- 1x Mirai Botnet or FireNet Botnet
- 50x Bots (FireNet)
- 500x Bots (Mirai)

Image:

<https://security.radware.com/ddos-threats-attacks/threat-advisories-attack-reports/mirai-rapid-evolution/>

More Addresses, More Business

Much more diverse data and capabilities available to attackers (eg. video, health stats, thermostat control).

Potential gains in being able to see important device information from certain types of IPv6 address, such as manufacturer and general level of security in the network.

Targeted, stealthier IPv6 scans could leave less of a trace compared to exhaustive IPv4 searches - easier to exclude irrelevant targets.



Image:

<https://www.mcafee.com/blogs/consumer/consumer-threat-notice/black-hat-danger-drones-thermostat-ransomware/>



Image: <http://www.bbc.co.uk/newsbeat/article/38824698/entrepreneurs-nick-jenkins-and-sarah-willingham-are-leaving-dragons-den>

Why Are You Doing This?

These tools **will** inevitably be developed by malicious actors as more and more hosts migrate to IPv6. There are already economic systems in place which will motivate this.

If we start working on possible approaches early, there is more time to develop strategies to defend against these scanning attacks (eg. heuristics to detect IPv6 address space reduction strategies in network traffic, recommendations to manufacturers).

Why Are You Doing This?

IPv6 is difficult to get right even in more complex operating systems.

IoT is one of the main drivers behind IPv6 adoption - raising security awareness about IPv6 among manufacturers is important, particularly where devices connect with each other.



 Candletouch Creator
4 months ago

Hi . Yes that is the plan to connect to Alexa and Google Home. Will still need you to be within sight of candle to light it.

 
4 months ago

Will we be able to connect this candle to Google Home/Amazon Alexa ?

Images: https://www.youtube.com/watch?v=5K1e_tWuwYU,
<https://www.kickstarter.com/projects/candletouch/candle-touch-the-first-smart-connected-real-flame-candle>

Research Question

Research Question

IPv6 addressing is said to be more secure than IPv4 due to it being impossible to exhaustively scan the 128-bit address space.

This is not necessarily true: a determined attacker can exploit IPv6 addressing conventions and publicly available information to locate hosts.

Embedded metadata in some IPv6 addresses can tell an observer a lot about the host it belongs to, and general level of security on the network.

IPv6 scans are potentially less noisy than exhaustive IPv4 scans.

Research Question

Would malware operators have a use for performing targeted IPv6 scans for specific types of devices instead of indiscriminate IPv4 scans?

Next Steps

Next Steps

See how many IoT devices are IPv6 capable, and attempt to locate them using a custom local network scanner.

This establishes whether it's actually possible to locate poorly secured IoT devices through IPv6 scans in the first place.




Image: <https://moniotrlab.ccis.neu.edu/wp-content/uploads/2018/09/lab.png>

Next Steps

Establish whether smartphones, laptops, and desktops, can be actively probed using MAC address-based IPv6 addresses.

This determines whether malware operators could locate specific types of complex devices through automated means, in addition to social engineering.



Image: https://manycore.org.uk/photos/photo_lab.jpg

Future Directions

Clean up the mess - exactly how can the findings from this initial project be used to prevent/mitigate these scanning attacks?

Issue recommendations to device manufacturers and possibly mainstream OS devs to fix identified risks.



Image:

<https://www.tradesmansaver.co.uk/wp-content/uploads/2019/06/commercial-cleaning-1920x570.jpg>

Malware and IPv6 Scanning

Vivian Band, University of Glasgow

Email: v.band.1@research.gla.ac.uk

Twitter: @ArcStatic42