

Advances in IPv6 Network Reconnaissance

Fernando Gont




LACNIC 29

Ciudad de Panamá, Panama. Abril 30 -Mayo 4, 2018

Introduction

Network Reconnaissance

reconnaissance

/rɪˈkɒnɪs(ə)ns/ 

noun

military observation of a region to locate an enemy or ascertain strategic features.

"an excellent aircraft for low-level reconnaissance"

synonyms: preliminary survey, [survey](#), [exploration](#), [observation](#), [investigation](#), [examination](#), [inspection](#), [probe](#), [scrutiny](#), [scan](#); [More](#)

- preliminary surveying or research.
"conducting client reconnaissance"

Network reconnaissance:

Locate possible targets and/or learn network information/features that can be leveraged for performing network-based attacks

Going mass scale

- What if we wanted to target the whole IPv6 Internet or a whole country?
- How do we find information about the “most popular” nodes?
- Some boring and dirty work needs to be done
 - What are the TLDs for a given region?
 - What are the suffixes for a given TLD?
 - etc

Going mass scale

- Some techniques need to be adapted or evaluated
 - e.g. dnsrevenue6 tend to fail on very short prefixes
- Other techniques need to be extrapolated
 - e.g. smarts on prefixes as opposed to addresses
- Where else to go and look for information?

Where to start?

Where to get to the most important bits?

- There were at least three datasets of popular sites:
 - Alexa's Top-1M Domains
 - Majestic's List
 - Umbrella list
- All available at: <https://github.com/fgont/domain-list>
- But far from the number of existing domain names...

Zone files for all

- Some TLDs zones (e.g. .ORG) shared via:

<https://czds.icann.org/>

- Some ccTLD zone made voluntarily available:

<https://zonedata.iis.se/>

- Some leaked:

<https://github.com/mandatoryprogrammer/RussiaDNSLeak>

Leveraging Search Engines

Challenges

- Most search engines support this sort of query:
site: *DOMAIN*
- Some engines obfuscate the results
 - Google is a notable example
- Some will ban you if they assume you are a robot
 - Teoma will ban you for about a day
- Some require you to keep state (cookie-like)
 - Just scrap the first page for the “cookie”, and use it in the actual query
- Some complain if they think you are a robot
 - Fake the user-agent
 - Fly low, if necessary

Playing with Teoma

- Good search results
- No obfuscation of results page
 - Improvements in scanning techniques
 - Improvements in IPv6 addressing to mitigate these attacks
- Banning upon lots of queries
 - Limits usefulness for a single target
- Example:

```
script6 get-teoma navy.mil
```

Playing with Bing

- Good search results
- No obfuscation of results page
 - Improvements in scanning techniques
 - Improvements in IPv6 addressing to mitigate these attacks
- No banning upon multiple queries
- Example:

```
script6 get-bing navy.mil
```

Playing with dictionaries

- Performance is much increased with the help of a dictionary
- Example:

```
script6 get-bing-dict navy.mil english.dic
```


Address patterns: Any changes?

Introduction

- Recent years saw publication of:
 - RFC7217
 - RFC8064
- Any changes?


Alexa Dataset

Interface Identifiers for web servers (Alexa)




Where to get to the most important bits?

Interface Identifiers for web servers (Umbrella)



Where to get to the most important bits?

Interface Identifiers for web servers (Majestic)



Conclusions

- Use of randomized increased to around 15%-20% for the worst-case scenario
- These figures didn't change much for mailservers or name servers
- Curiosity: there **was not** a move from IEEE-based -> randomized

Notes on DNS reverse mappings

Introduction

- DNS reverse mapping is among the most powerful techniques for IPv6 enumeration
- We learned some lessons...

“Noise”

- Large number of dynamically generated reverse mappings for some networks:

```
Found: 2001:4998:c:80d::4062 is hz-network-migration-50568-89.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4064 is hz-network-migration-50568-91.gq1.yahoo.com.  
Found: 2001:4998:c:80d::406d is hz-network-migration-50568-100.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4061 is hz-network-migration-50568-88.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4066 is hz-network-migration-50568-93.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4060 is hz-network-migration-50568-87.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4063 is hz-network-migration-50568-90.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4068 is hz-network-migration-50568-95.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4069 is hz-network-migration-50568-96.gq1.yahoo.com.  
Found: 2001:4998:c:80d::406b is hz-network-migration-50568-98.gq1.yahoo.com.  
Found: 2001:4998:c:80d::4065 is hz-network-migration-50568-92.gq1.yahoo.com.  
Found: 2001:4998:c:80d::406f is hz-network-migration-50568-102.gq1.yahoo.com.  
Found: 2001:4998:c:80d::406c is hz-network-migration-50568-99.gq1.yahoo.com.
```

Reliability

- Reverse mappings of /48s were more reliable than those of /32s
- May make sense to split /32s into multiple /48s for reliability purposes

Integrating IPv6 Network Reconnaissance

Introduction

- Most network reconnaissance is manual
- Our goal was to try to integrate different techniques into the same tool

Messi: IPv6 net reconnaissance tool

- If you have access to a local node, it might be of use:
- What the tool does:
 - 1) Obtain domains from search engines
 - 2) Obtain NS and MX records
 - 3) Obtain IPv6 addresses for all those names
 - 4) Build prefixes out of those addresses
 - 5) Do DNS reverse enumeration
 - 6) Go back to step #1
- Eventually we converge to results

Questions?

Thanks!

Fernando Gont

fgont@si6networks.com

IPv6 Hackers mailing-list

<http://www.si6networks.com/community/>



www.si6networks.com