Understanding the Implementation and Security Implications of Protective DNS Services

Paper Presentation

Patrick Mao

Problems with DNS

- Botnet command and control (C&C), phishing, spam, and malware distribution
- 91% of Internet attacks are from resolving malicious domain names (Cisco)
- In March 2023, DAAR reported over 622k malicious domains

Current Countermeasures

- Domain takedown: cumbersome procedure
- Protective DNS (PDNS)

Protective DNS

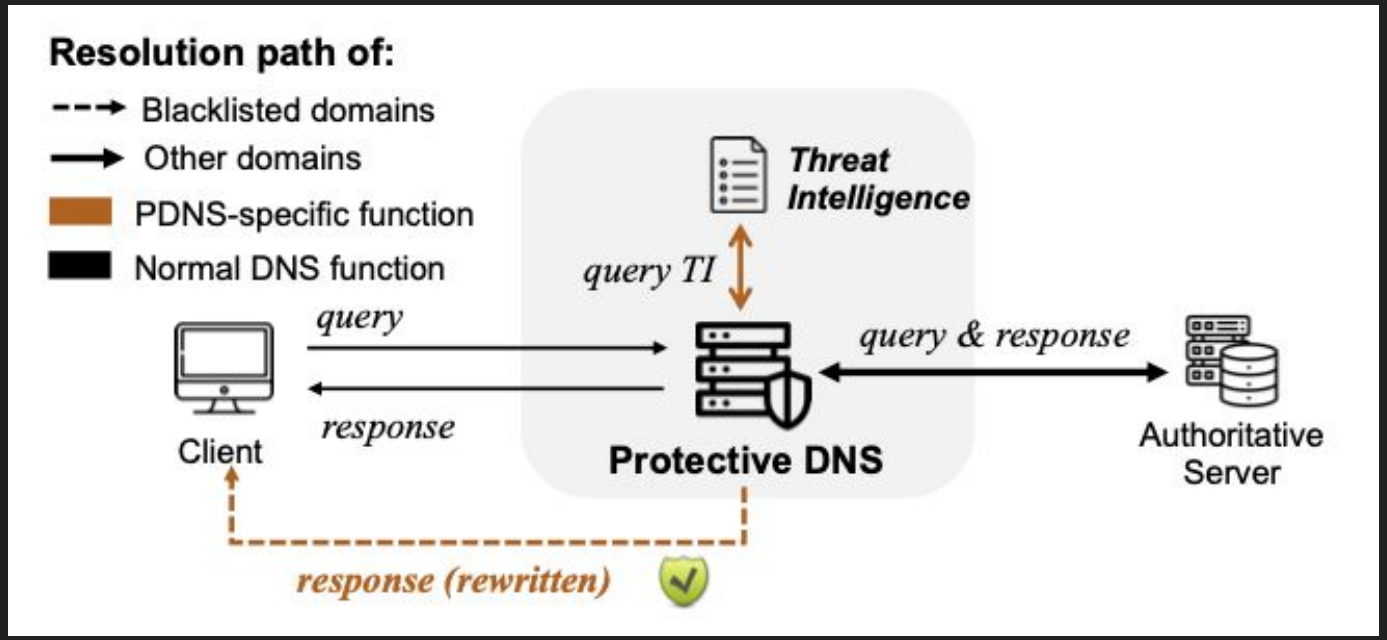

Research Questions

- How many DNS servers provide PDNS?
- What are the blocking policies?
- Any security risks?

How many DNS servers provide PDNS?

Straightforward because all PDNSes block through response rewriting

1. Collect malicious domain names
2. Query target DNS servers and authoritative DNS servers (3 vantage points)
3. Compare responses

Challenges in Identifying PDNSes

Distinguishing modified responses from other DNS manipulations

- Determine whether the response is rewritten
  - Studies show if the DNS response IPs do not share ASN with auth servers, it's likely rewritten (thoughts?)
- Exclude Censorship Induced Rewriting
  - Query from countries with high internet freedom (US, UK, JP)
  - Don’t include potentially censored domains e.g. political sites
  - Sending test domains to a random IP in the target AS that's not a DNS resolver. If it returns an answer, then it must be injected by a censor. (issues?)
- Exclude DNS Hijacking Induced Rewriting
  - Distributed queries

Final Trick: Only consider a resolver PDNS if it rewrites > threshold number of answers

The PDNS Scanning System

1. Collect malicious domain names
    a. Collected 36K highly malicious domains
    b. Randomly sampled 10K and use as the final list
    c. Use Tranco 100 as the control group to test DNS availability
2. Query target DNS servers and authoritative DNS servers (3 vantage points)
    a. Only target DNS servers are stable over a month (193K stable resolvers)
    b. Use XMap to query all selected DNS servers for all 10,100 domains
    c. 30 rounds of querying, each round with 10,100 domains, log rewritten responses
3. Compare responses

Measurements

- 17K PDNSes identified from 193k stable resolvers
  - 9% adoption rate
- US has 21% adoption rate, China has 4.5%
- User-side Adoption (from Netflow datasets)
  - 9,470,810 DNS queries analyzed (~25K per day)
  - 24K out of 33K unique clients use PDNS (73% 😮)
  - ~800 PDNS queries per day from a single Chinese college campus (issues?)

CC	# IP	ASN	# IP
US	6,296 (35.8%)	20115 (CHARTER-20115)	1,074 (6.1%)
IRN	1,225 (7.0%)	3303 (SWISSCOM)	777 (4.4%)
CN	1,205 (6.8%)	209 (CenturyLink Communications)	705 (4.0%)
JP	1,056 (6.0%)	5617 (TPNET)	613 (3.5%)
CH	804 (4.6%)	17506 (UCOM)	576 (3.3%)
PL	745 (4.2%)	10796 (TWC-10796-MIDWEST)	570 (3.2%)
MD	635 (3.6%)	21342 (AKAMAI-ASN2)	523 (3.0%)
ID	540 (3.1%)	8926 (MOLDTELECOM-AS)	480 (2.7%)
OM	380 (2.2%)	2519 (VECTANT)	420 (2.4%)
RO	367 (2.1%)	50010 (Nawras-AS)	379 (2.2%)
117 Countries		1,473 ASNs	

Measurements

Querying Performance

- PDNS does not incur performance overhead

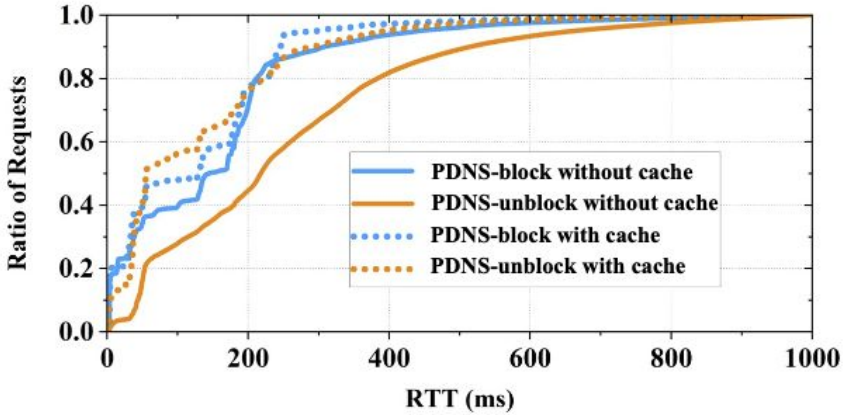

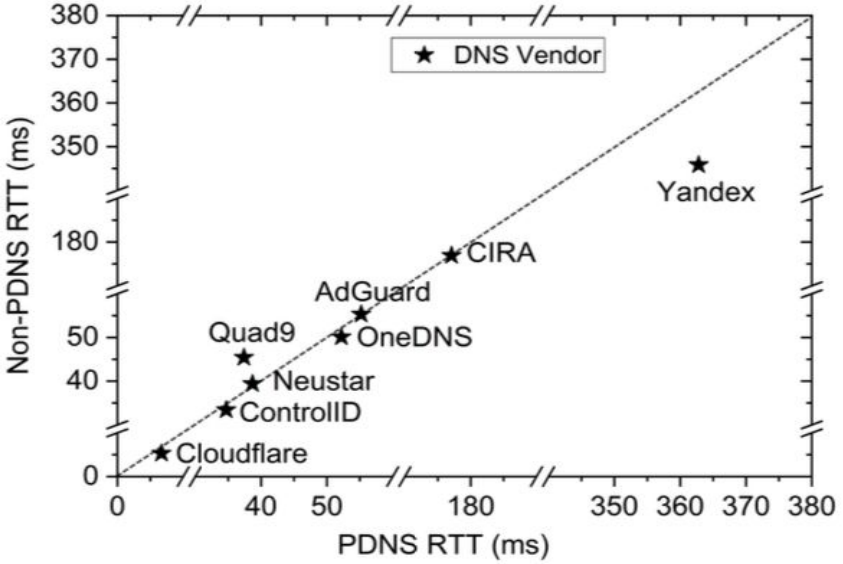

Measurements

Blocked domains

- PDNSes tend to only block high-risk domains

TABLE VI: Category of domains blocked by PDNSes.

Category	# Test domains	# Avg. blocked domains	PDNS Coverage
Malware	4,231	961.9	17,596 (99.97%)
Botnet	3,962	472.0	17,529 (99.59%)
Phishing	867	160.9	17,213 (97.80%)
Adult	667	119.8	12,680 (72.04%)
Spam	259	96.6	16,628 (94.47%)
Tracker	14	0.5	3,779 (21.47%)

Measurements

Similarities map

- Black spots indicates high similarity
- Mostly not similar. Why?
- Grouping might suggest inability in distinguishing PDNS rewrites from censorship rewrites.

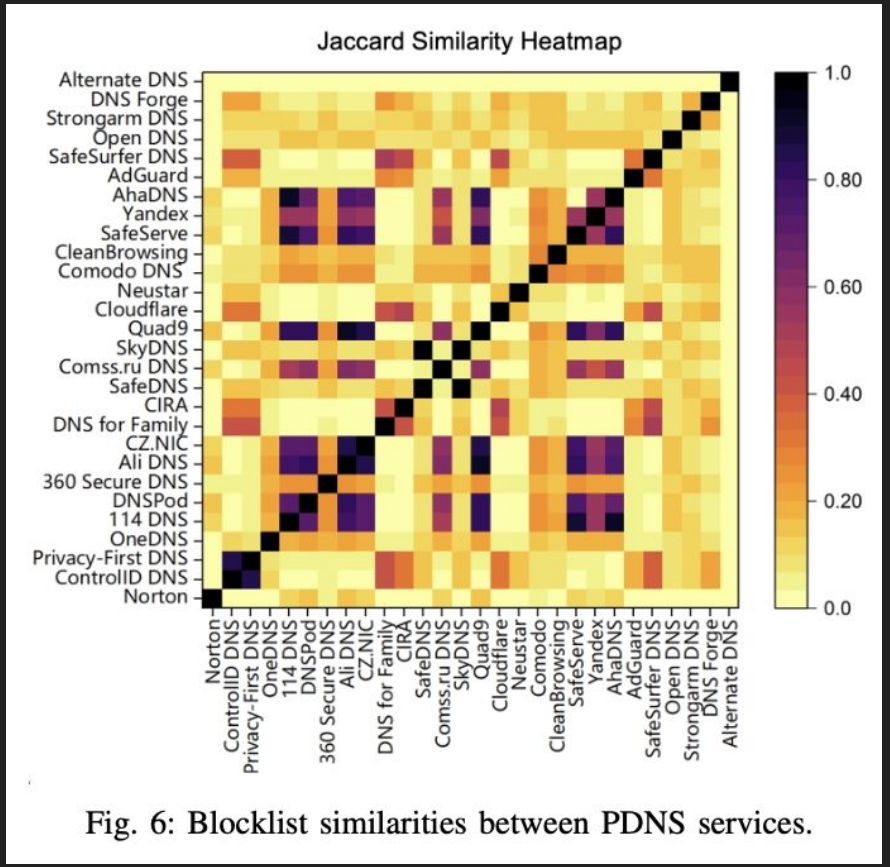

Fig. 6: Blocklist similarities between PDNS services.

Security Issues

Total number of PDNSes = 17K

Denial of Response

- 28 PDNS servers block the source IPs that query malicious domains
- Attacker can simply query the PDNS with victim's IP
- The authors tried this on all 28 PDNS servers and all of them blocks the researchers’ source IP

Dangling PDNS Infrastructure

- 26 PDNS servers return addresses pointing to dangling cloud infrastructure
- Domain/IP takeover

More Security Issues

Total number of PDNSes = 17K

Flawed (Loose) implementation of PDNS

- 105 PDNSes returned both rewritten answers and authoritative answers for malicious domain queries
- PDNS operators probably do this to mitigate risk of complete disablement of (erroneously) blocked domains

Non-configured query types of PDNS

- 13 PDNSes return the original resolution results for types are not configured (e.g. TXT records)

Mitigation Recommendations

Transparent blocking activity

- To boost user experience, set up a webpage to inform user that the site they are visiting is harmful. and providing channels for complaints

Utilizing safe rewriting infrastructures

- TLDR; don't be lazy and use third-party website as a rewrite target

Defense of denial of response

- No clean answer to this because DoR does serve valid purposes such as preventing botnets talking to their C&C
- The authors stated one way: if a client issues numerous request to malicious domains, reply with a large DNS answer to force it to use DNS over TCP. (how is this effective?)

The Class’s Ratings

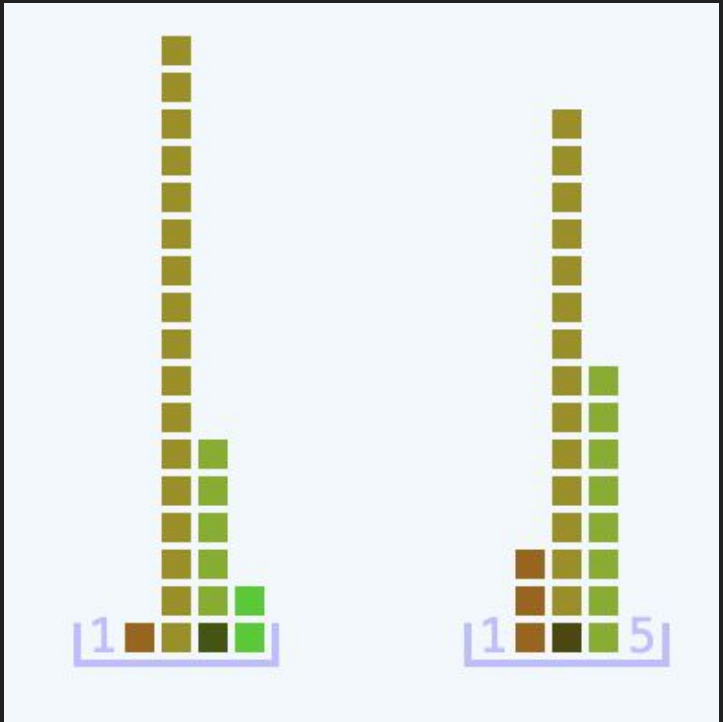

- The “Most 3s” award
- Quality: 12 in all 16 and last in Network
- Interest: 14 in all 16 and last in Network

What has the paper done well

1. **Comprehensive Measurement**: The first large-scale measurement of the PDNS ecosystem, identifying 17,600+ PDNS servers and analyzing their adoption trends.
2. **Practical Impact**: Uncovers and validates multiple security vulnerabilities in PDNS in the wild.
3. **Methodology:** Scalable methodology for identifying PDNS services

Improvements and Next Steps

1. **Limitations of PDNS blocklists**: Open-source lists are often incomplete;
2. **Explore Other Attacks/Defenses**: Residential DNS resolvers; more PDNS attacks; understanding the severity; real-time PDNS adaptations; monitoring tools for malicious rewrites
3. **PDNS standardization**: Consensus on different countries' PDNS guidelines; Study why implementations and blocklist adoption diverge (and converge for some)

Discussions

- PDNS is good to prevent malware but what if governments use it for censorship?
- How are things adopted so easily without any heavy testing in this domain?
- How do they update blocklists?
- Is it possible to exploit PDNS to attack availability of normal websites?
- Do PDNS needs a standard? How do we establish a standard for PDNS?