UNIVERSITÄT DUISBURG ESSEN

Internet Measurement Research

Matthäus Wander

<matthaeus.wander@uni-due.de>

Kassel, October 1, 2013

Overview

- How to get measurement data?
- Research projects
  - Case studies of past projects
  - Ideas and inspiration for new projects
- Measurement methods, challenges, solutions

How to Get Measurement Data?

- Existing data collected by other researchers
  - Research papers are on the web, but not the data
- Obstacles:
  - Privacy issues, fear of abuse
  - Documentation and anonymization effort
  - Collected data is often bound to one purpose
- www.DatCat.org measurement data catalog
- Semi-public data: find operator, sign NDA
- Collect your own

Collecting Data

- Passive: monitor existing traffic
  - Test your own network? Persuade NOC?
- Active: probe networks and hosts
  - Effort vs. data quality (time/bandwidth/latency/loss)
  - Just active sender or also active receiver?
  - Prepare for complaints with active probing
- Save raw data if possible
  - You may want to further analyze unexpected effects

Research Projects

1. Determine behavior of NAT routers ...6
2. Count DNSSEC validating clients ...9
3. Global impact of DNS censorship ...17
4. Analysis of public DNSSEC keys ...30
5. Effectiveness of DNS caches ...32
6. Analysis of darknet traffic ...35

Determine Behavior of NAT Routers

Research Project 1

Measurement Method

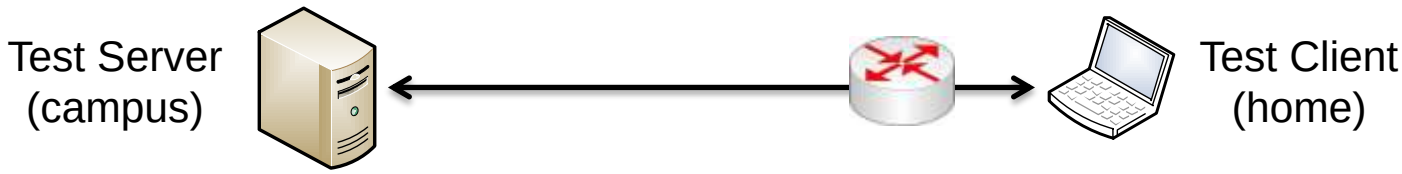

- Active measurement between two test programs
  - Client in user home network
- User must download+run Windows tool
  - Required the then new .NET 4 (no Linux/Mac version)
  - Tool needs raw socket and WinPcap (admin privileges)
- Incentive for users:
  - „help us for science“

Measurement Method (2)

- Asked students and friends to run the tool
  - 40 usable results in 2 weeks
- Multiple tests, each repeated 3 times
  - Found some anomalies by repeating same test
- Send result for each test to our server
- Manual result analysis with Excel
- Some results suggested more detailed analysis
  - Which wasn‘t possible, raw IP packets not saved

Count DNSSEC Validating Clients

Research Project 2

Measurement Method

- How many web clients are protected by DNSSEC?

```
<img src="http://dnssec.vs.uni-due.de/r/a" alt="" height="1" width="1">
<img src="http://dnssec.vs.uni-due.de/r/b" alt="" height="1" width="1">
```

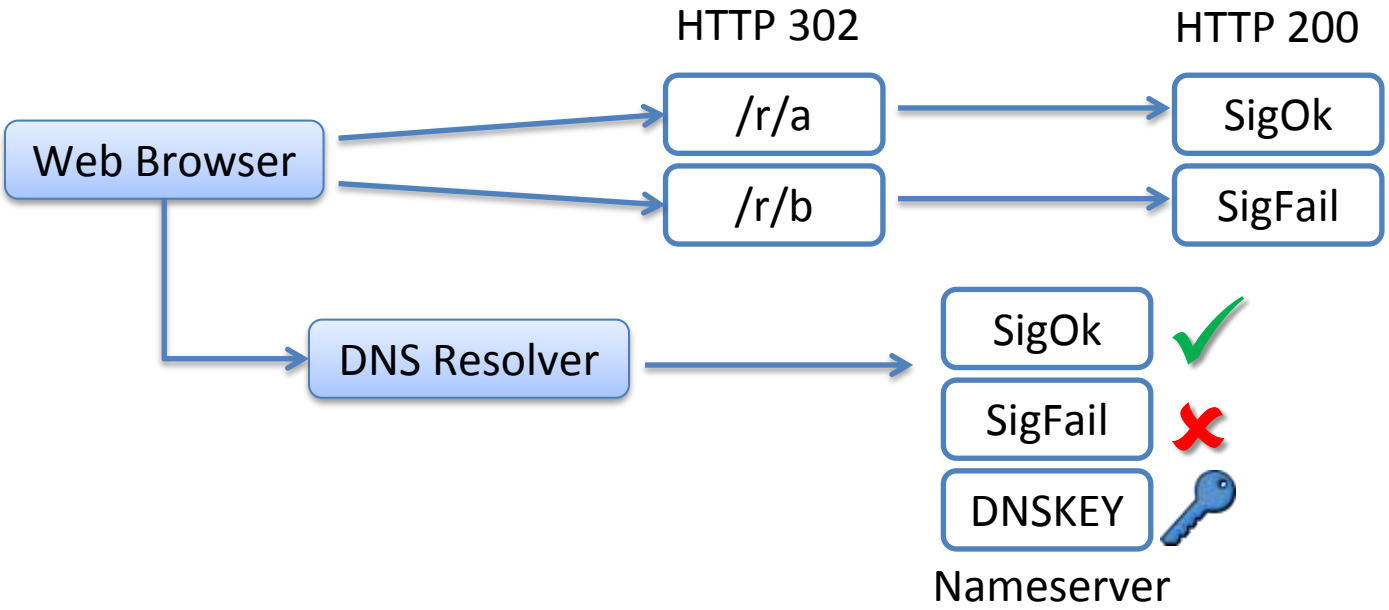

- How to generate page impressions?

Embed HTML Snippet in Popular Websites

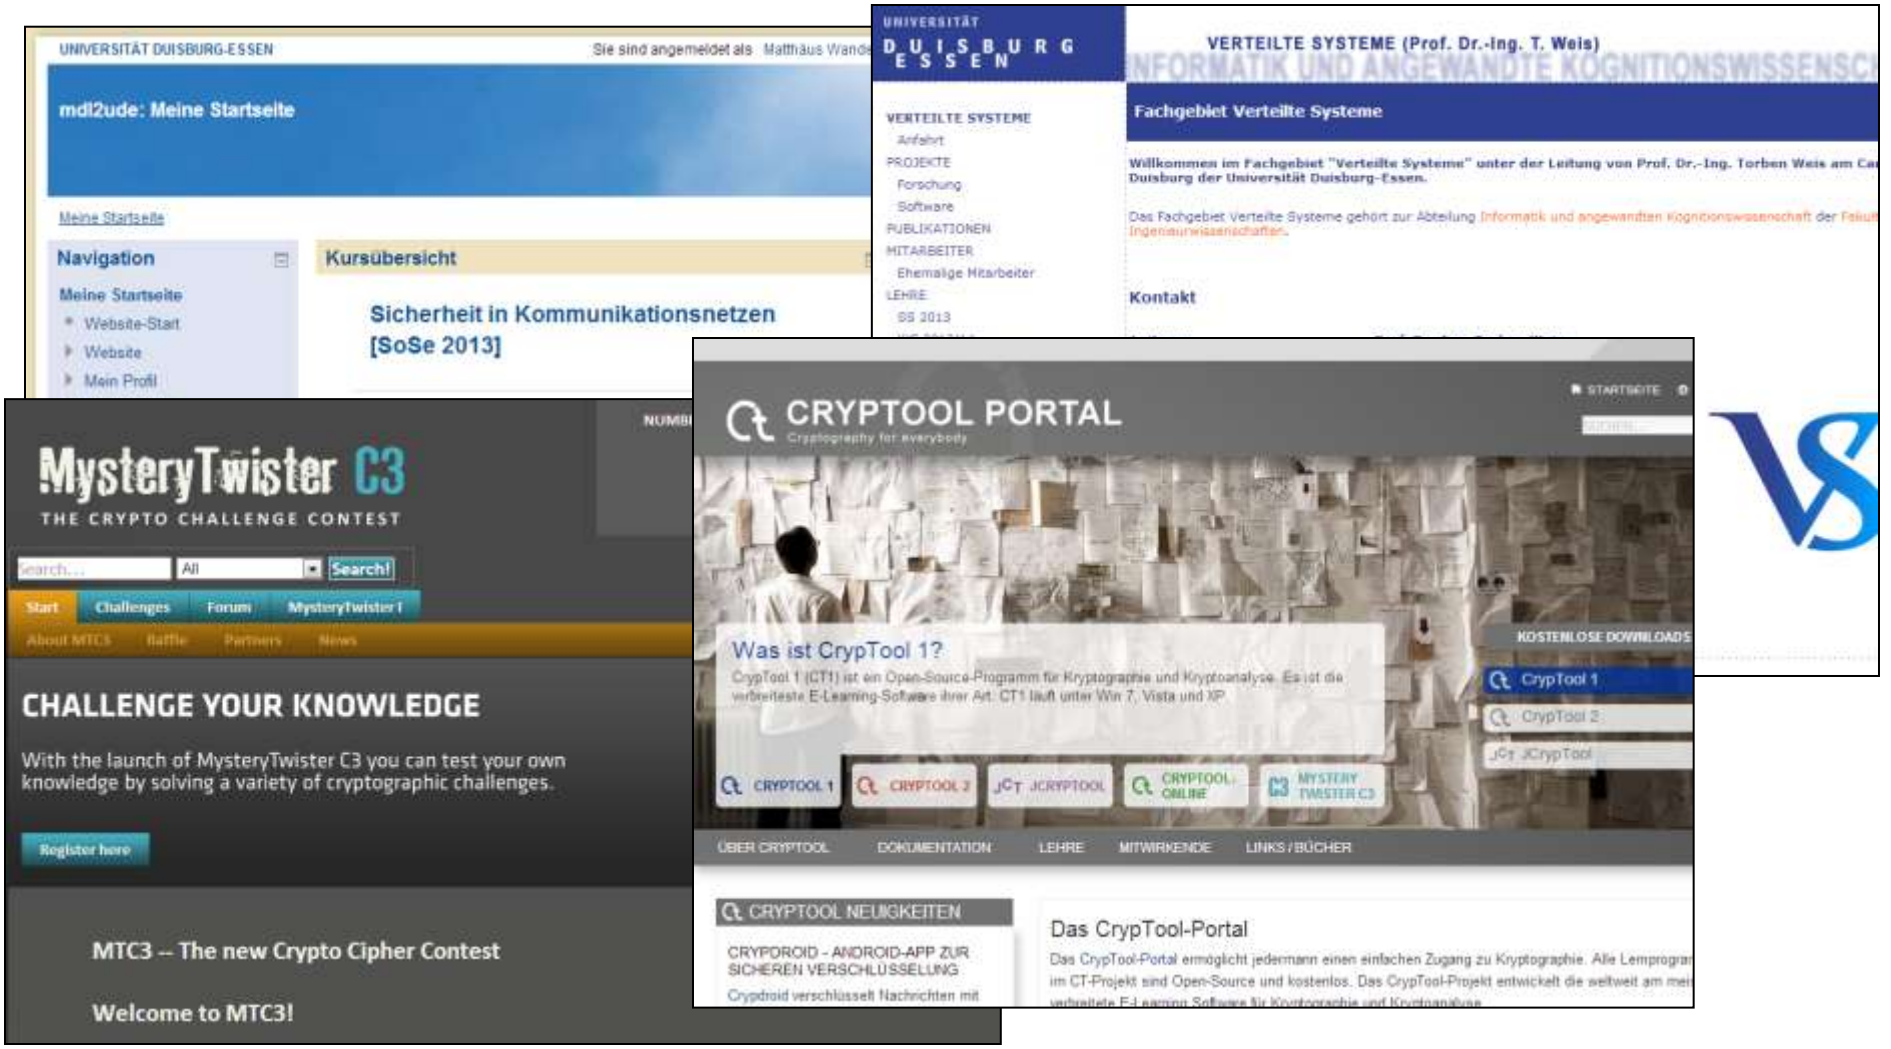

Autosurf Communities

- Sign up with „autosurf“ traffic exchanges
- Automated website visits from various clients

Webpage with Active Measurement

DNSSEC Resolver Test

This test determines whether your DNS resolver validates DNSSEC signatures. For this test you need JavaScript turned on.

Start test

Most people will experience a negative test result (no DNSS

Help Us

Point your friends to this webpage to help us measure the spr

If you are operating a website and would like to help us, con

DNSSEC for Users

Few operating systems support DNSSEC validation out of t
information). Keep in mind that web browsers do not distingu

Google dnssec test

Web Images Maps Shopping Applications More Search tools

About 1,160,000 results (0.11 seconds)

Ad related to **dnssec test**

DNSSEC Tools - Set Up Your Own DNS Server - InternetSociety.org
www.internetsociety.org/**DNSSEC**-Tools
DNSSEC Testing Tools and Resources
Internet Society Deploy360 Programme has 237 followers on Google+
$10,000 Community Grants - Apply for Grants & Awards - Enable IPv6

DNSSEC Analyzer
dnssec-debugger.verisignlabs.com/
The **DNSSEC** Analyzer from VeriSign Labs is an on-line tool to assist with diagnosing problems with **DNSSEC**-signed ... Enter a domain name to be **tested**.

DNSSEC Resolver **Test**
dnssec.vs.uni-due.de/
DNSSEC Resolver **Test**. This **test** determines whether your DNS resolver validates **DNSSEC** signatures. For this **test** you need JavaScript turned on. derp.

No, you are not using **DNSSEC**
dnssec-or-not.org/
2011-2012 VeriSign, Inc. All rights reserved. VERISIGN, the VERISIGN logo, and other trademarks, service marks, and designs are registered or unregistered ...

Domaininfo - **DNSSEC** Checker - **Test** your DNS
domaininfo.**dnssectest**.nu/
With our **DNSSEC test**, you can check that everything is in order with your domain, whether you have DNSSEC as a complement to your DNS, and that the ...

Approaches in Related Projects

- Geoff Huston (APNIC):
  - Buy 350.000 web hits with $100 Flash advertisement
  - Use Flash to query DNSSEC-signed domain names

```
GET http://t10000.u5950826831.s1347594696.i767.v6022.d.t5.dotnxdomain.net/1x1.png
GET http://t10000.u5950826831.s1347594696.i767.v6022.e.t6.dotnxdomain.net/1x1.png
```

- Duane Wessels (Verisign):
  - Register wpad.$tld *(RFC 3040 Web Proxy Autodiscovery Protocol)*
  - Use DNS-only technique to identify DNSSEC validators

Result Analysis

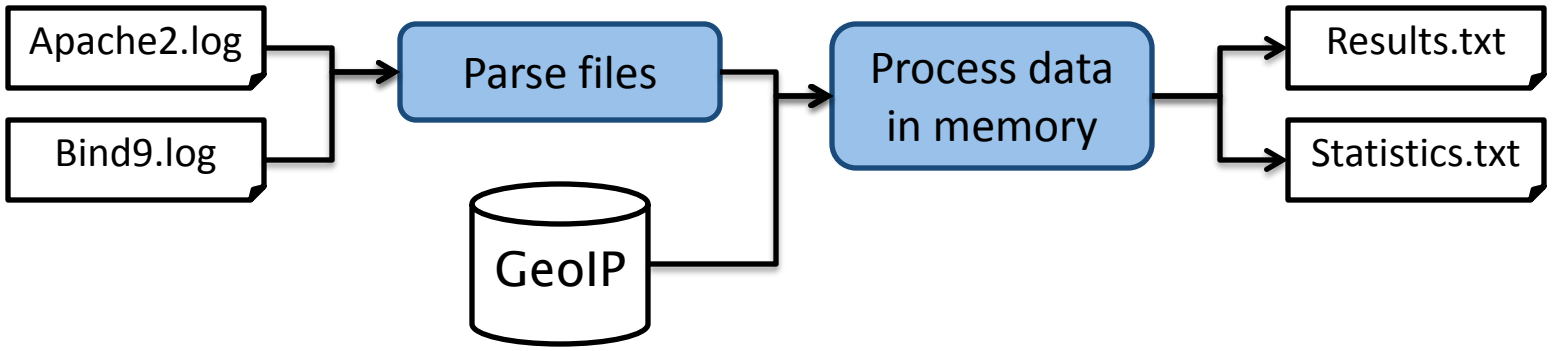

- Parse whole files into memory, then analyze
  - Does not scale with large log files
- Pipeline parsing and processing
  - Still needs to parse all log files (>10 GB)
  - Ideal: incremental analysis, results on website

Selection Bias

- Is this result representative?
- Group results by country code, AS number, etc.

AS	$\frac{V}{V_{total}}$	$\frac{V}{V+N}$	cli=dns
Comcast 7922	29.1%	69.0%	0.5%
KabelBW 29562	14.3%	86.4%	0.3%
M-Net 8767	6.1%	46.6%	3.9%
Telia SE 3301	3.3%	73.8%	1.5%
O2 CZ 5610	3.0%	69.2%	0.5%

UNIVERSITÄT DUISBURG ESSEN

Global Impact of DNS Censorship

Research Project 3

Scenario and Objectives

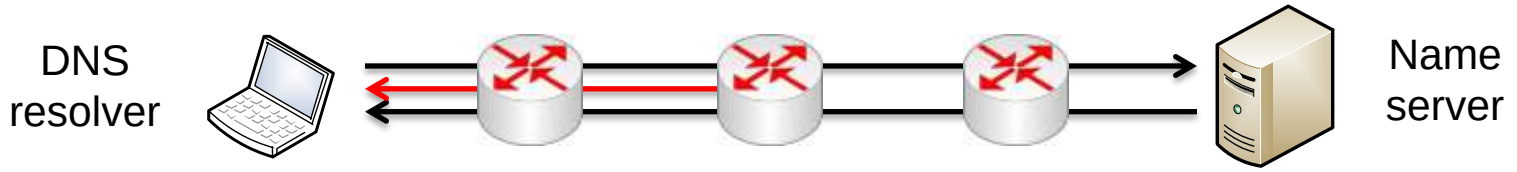

- DNS injection: IP router spoofs DNS response
- Send DNS requests to random IPv4 addresses
  - Check who responds with spoofed answer

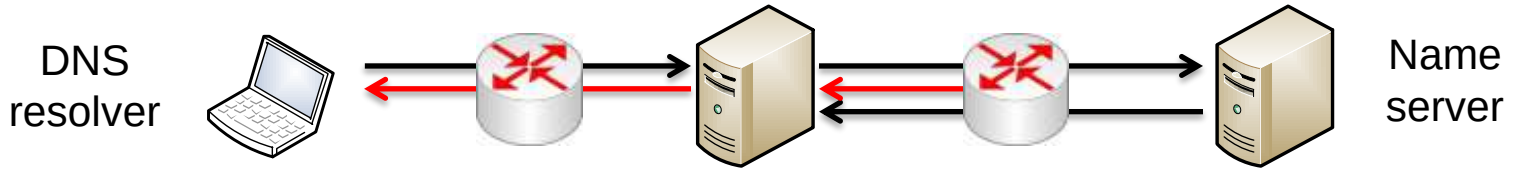

- Send DNS request to open resolvers worldwide
  - Check who is affected by DNS injection

Prepare for Complaints

- Announced measurement to our NOC
  - Prepared mail response template
- Ideal: get AS number to receive abuse mails
- Set up rDNS name and website on scanner host
  - crawler.vs.uni-due.de
- Contact information
- Offer blacklisting
- State purpose of scan

Probing the IPv4 Internet

- a.b.c.99: one query into each /24 subnet
  - Omitted 0/8, 10/8, 172.16/12, 192.168/16, 224/3
  - Ideal: omit BGP prefixes not globally announced

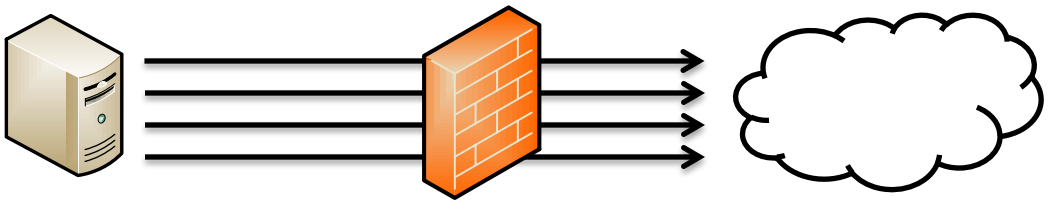

- Firewall creates state for each query packet
  - Extra CPU cost and limited space (Linux: 2^{16} entries)
  - Drops responses with wrong source port/IP address
  - Use packet filter rules without stateful inspection

Sending Queries and Monitoring Packet Loss

- DNS queries are UDP datagrams
  - One socket+port suffices for all 14M queries
- How to differ *no service* from *packet loss*?
  - ICMP errors might be an indicator (unused here)
- Monitor network load with periodic DNS pings
  - Two responders to identify origin of packet loss

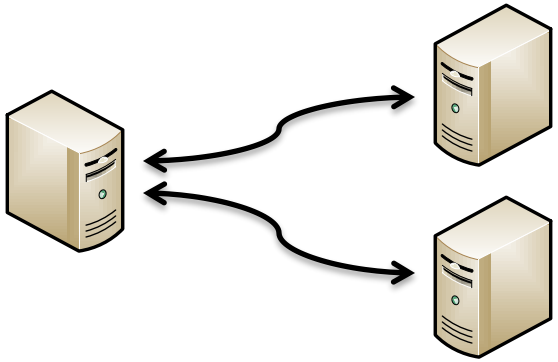

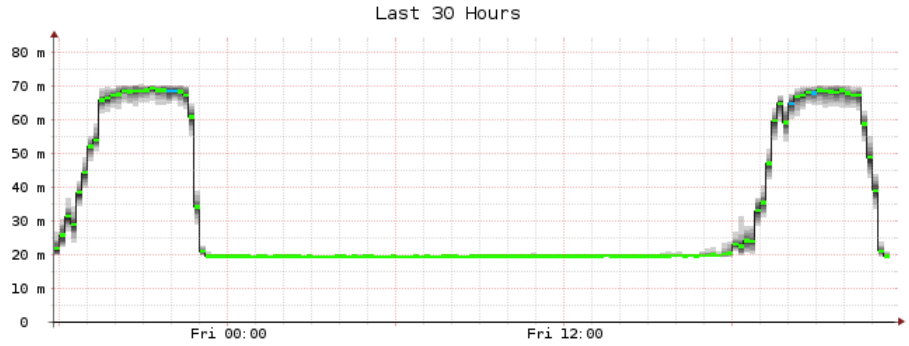

Avoiding Packet Loss

- First version sent ~22k queries/s
  - Problem: the faster it ran, the less responses arrived
- Naive solution: wait() to limit sending rate
  - Problem: 200 q/s (~20 KB/s) killed campus router
- Spread load per destination network over time

```
for c in range(256):
  for b in range(256):
    for a in range(256):
      yield "{0}.{1}.{2}.99".format(a, b, c)
```

Receiving Responses

- Ensure socket receive buffer doesn‘t overfill
  - Increase buffer size (SO_RCVBUF)
  - Read from socket in tight recv() loop
- Output: SQLite db without indices (fast write)
  - Saved parsed responses (less disk, less information)
  - Evaluation: recreated database with indices
- Receiving data with socket vs. packet capture
  - Socket misses faulty UDP checksums, ICMP errors
  - Pcap misses packets under high load

Results: Responses for “facebook.com”

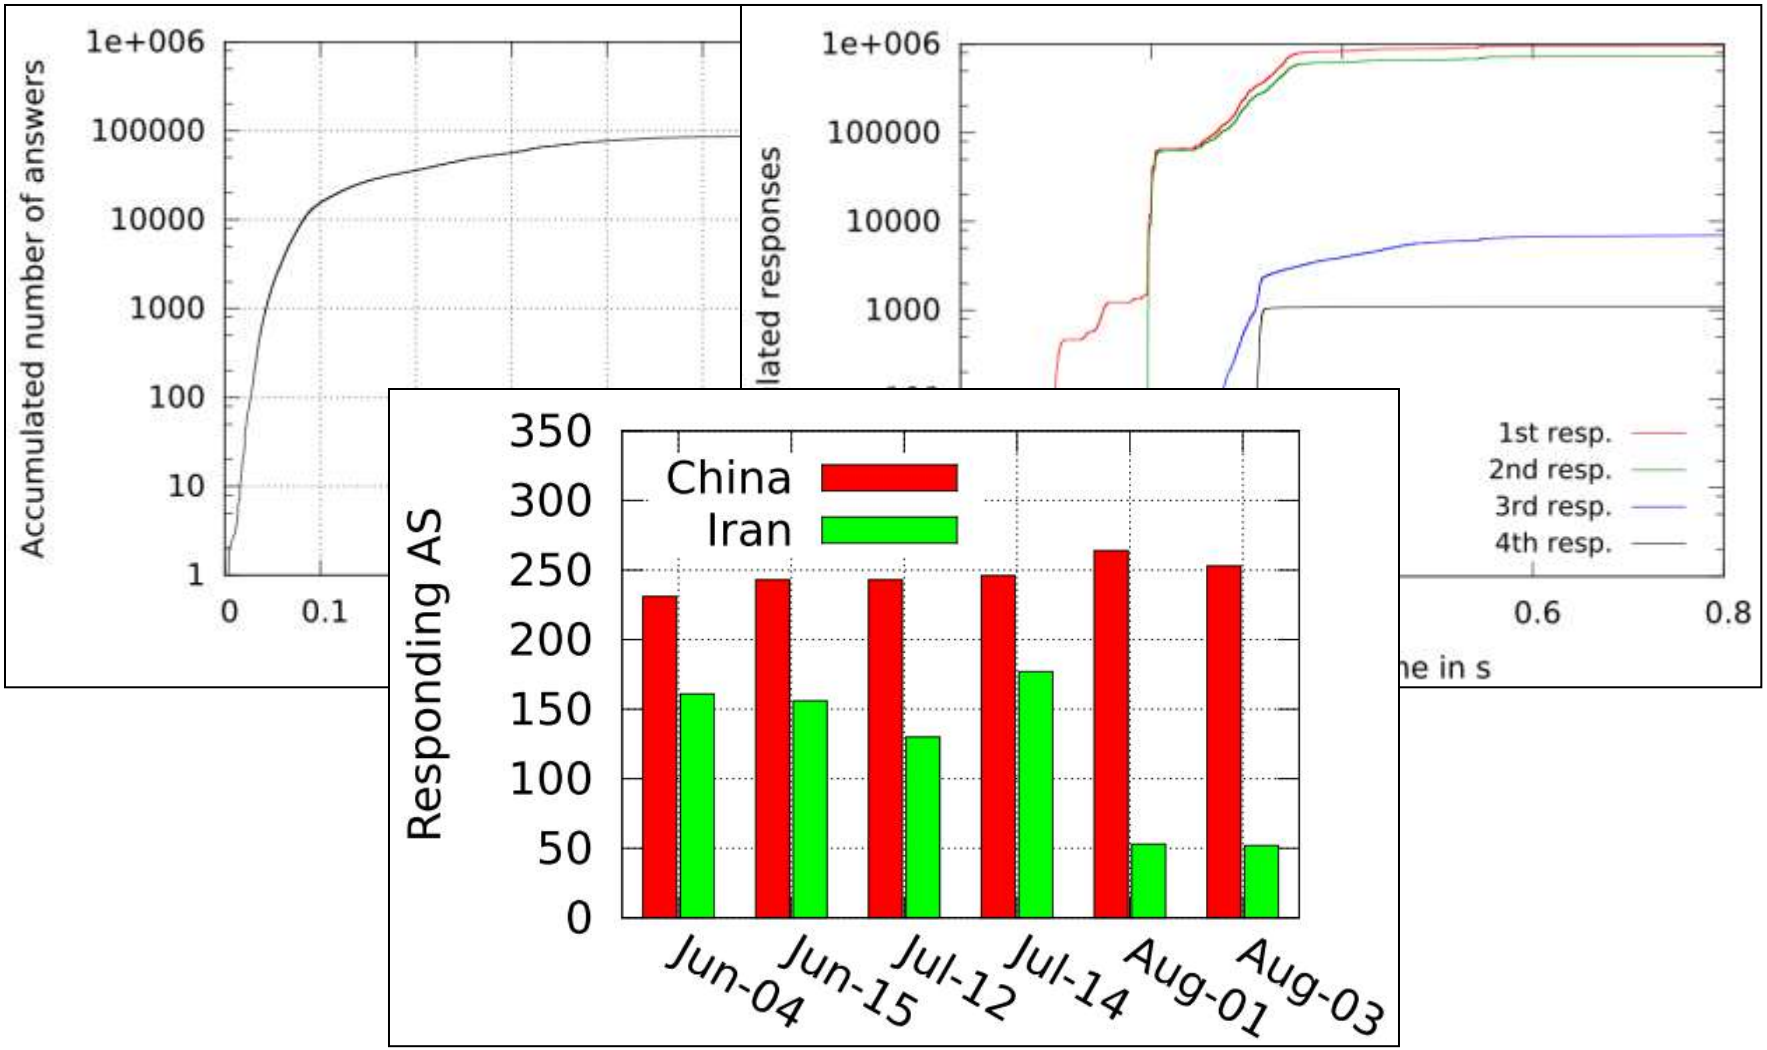

Querying Open Resolvers Worldwide

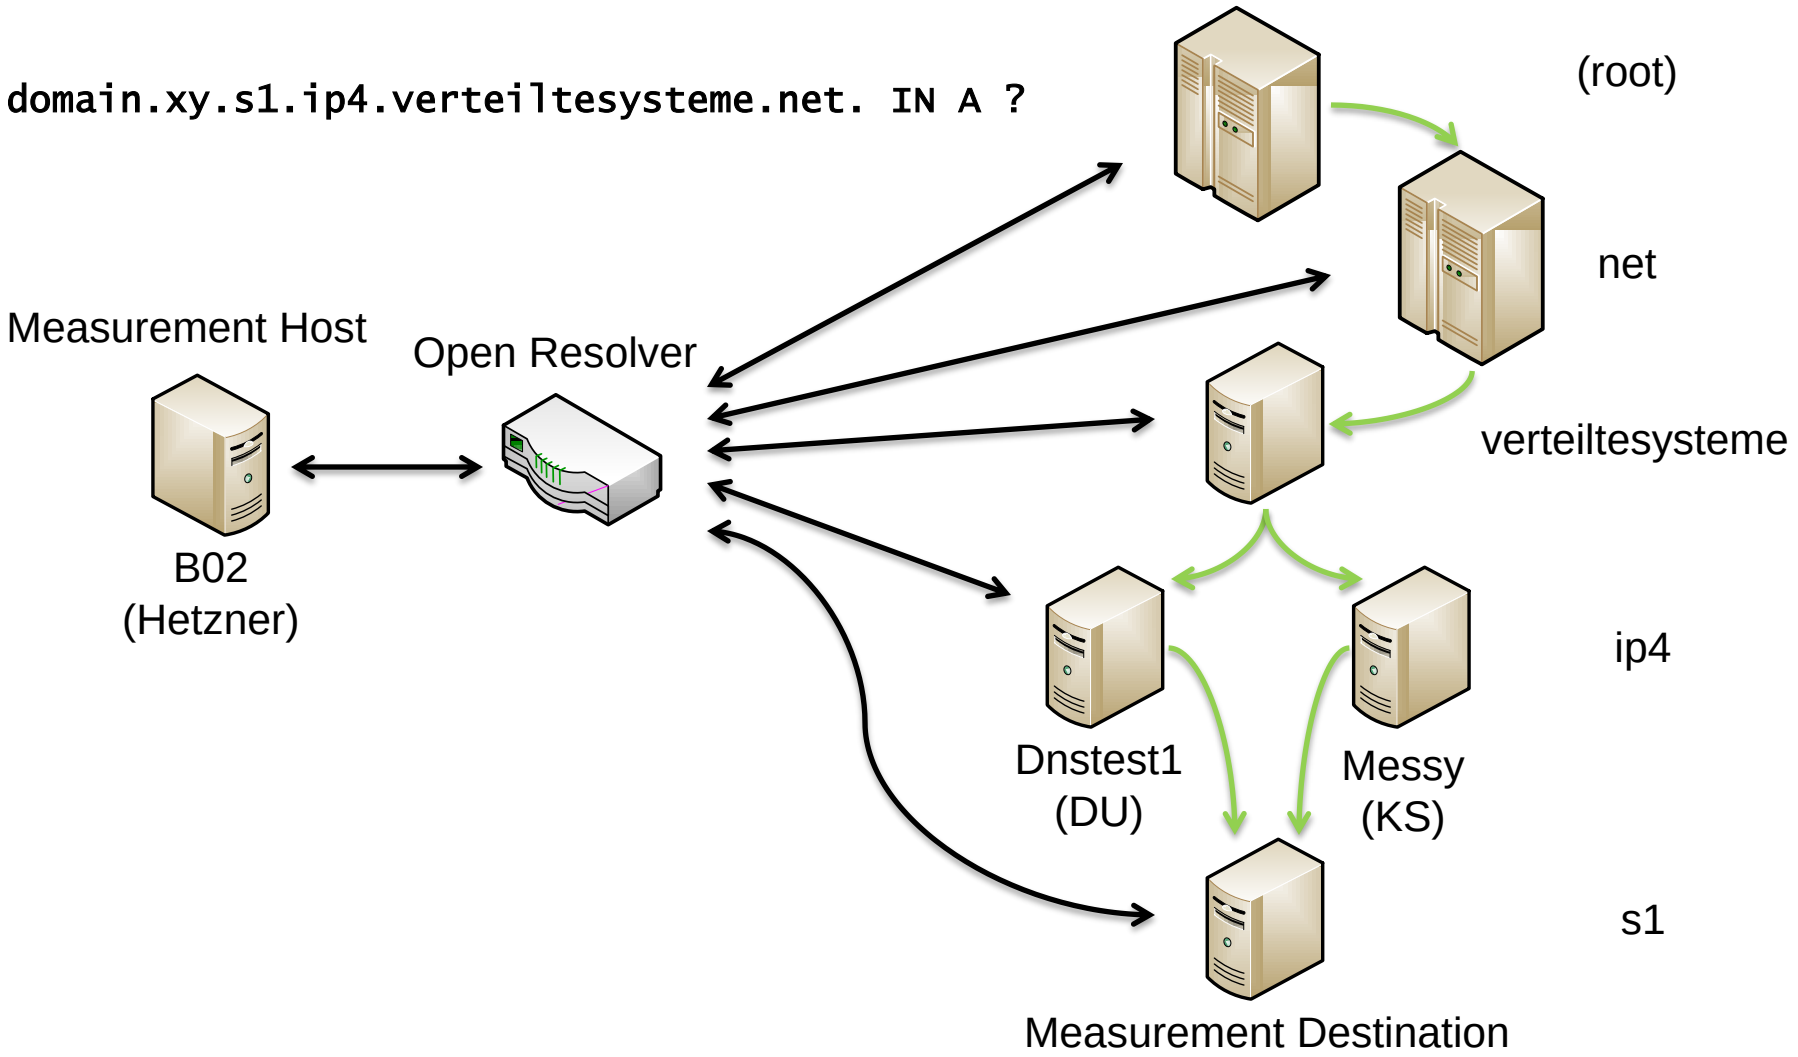

Open Resolvers and Measurement Destinations

- OpenResolverProject.org: 1M (out of 25M)
- root.zone: 1155 root and TLD name servers

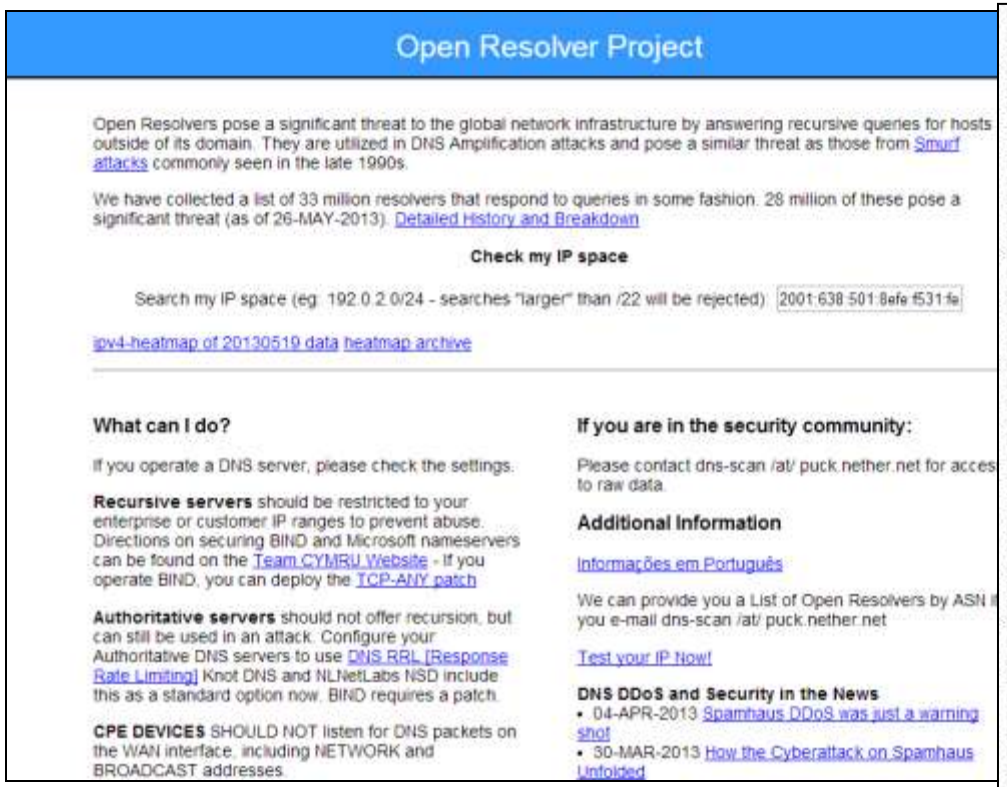

Open Resolver Project

Open Resolvers pose a significant threat to the global network infrastructure by answering recursive queries for hosts outside of its domain. They are utilized in DNS Amplification attacks and pose a similar threat as those from Smurf attacks commonly seen in the late 1990s.

We have collected a list of 33 million resolvers that respond to queries in some fashion. 28 million of these pose a significant threat (as of 26-MAY-2013). Detailed History and Breakdown

Check my IP space

Search my IP space (eg: 192.0.2.0/24 - searches "larger" than /22 will be rejected): 2001:638:501:8efe:f531:fe

ipv4-heatmap of 20130519 data heatmap archive

What can I do?

If you operate a DNS server, please check the settings.

Recursive servers should be restricted to your enterprise or customer IP ranges to prevent abuse. Directions on securing BIND and Microsoft nameservers can be found on the Team CYMRU Website - If you operate BIND, you can deploy the TCP-ANY patch

Authoritative servers should not offer recursion, but can still be used in an attack. Configure your Authoritative DNS servers to use DNS RRL [Response Rate Limiting] Knot DNS and NLNetLabs NSD include this as a standard option now. BIND requires a patch.

CPE DEVICES SHOULD NOT listen for DNS packets on the WAN interface, including NETWORK and BROADCAST addresses.

If you are in the security community:

Please contact dns-scan /at/ puck.nether.net for acces to raw data.

Additional Information

Informações em Português

We can provide you a List of Open Resolvers by ASN you e-mail dns-scan /at/ puck.nether.net

Test your IP Now!

DNS DDoS and Security in the News

- 04-APR-2013 Spamhaus DDoS was just a warming shot
- 30-MAR-2013 How the Cyberattack on Spamhaus Unfolded

```
bi.                     172800  IN      NS      bi.cctld.authdns.ripe.net.
bi.                     172800  IN      NS      ns.nic.bi.
bi.                     172800  IN      NS      dns.princeton.edu.
bi.                     172800  IN      NS      ns1.nic.bi.
bi.                     172800  IN      NS      anyns.nic.bi.
bi.                     172800  IN      NS      ns-bi.afrinic.net.
anyns.nic.bi.           172800  IN      A       204.61.216.61
anyns.nic.bi.           172800  IN      AAAA    2001:500:14:6061:ad:0:0:1
ns.nic.bi.              172800  IN      A       196.2.8.205
ns1.nic.bi.             172800  IN      A       196.2.12.205
bi.                     86400   IN      NSEC    biz. NS RRSIG NSEC
bi.                     86400   IN      RRSIG   NSEC 8 1 86400 20130724000000 201307
biz.                    172800  IN      NS      a.gtld.biz.
biz.                    172800  IN      NS      b.gtld.biz.
biz.                    172800  IN      NS      c.gtld.biz.
biz.                    172800  IN      NS      e.gtld.biz.
biz.                    172800  IN      NS      f.gtld.biz.
biz.                    172800  IN      NS      k.gtld.biz.
BIZ.                    86400   IN      DS      21910 8 1 5EAA597F7A5D92ECB6D8628044
BIZ.                    86400   IN      DS      21910 8 2 7C3B5FF5E65827A3D7CE2394B6
BIZ.                    86400   IN      RRSIG   DS 8 1 86400 20130724000000 20130716
a.gtld.biz.             172800  IN      A       156.154.124.65
a.gtld.biz.             172800  IN      AAAA    2001:503:7bbb:FFFF:FFFF:FFFF:FFFF:FF
b.gtld.biz.             172800  IN      A       156.154.125.65
c.gtld.biz.             172800  IN      A       156.154.127.65
e.gtld.biz.             172800  IN      A       156.154.126.65
f.gtld.biz.             172800  IN      A       209.173.58.66
f.gtld.biz.             172800  IN      AAAA    2001:500:3682:0:0:0:0:12
k.gtld.biz.             172800  IN      A       156.154.128.65
k.gtld.biz.             172800  IN      AAAA    2001:503:e239:0:0:0:3:2
biz.                    86400   IN      NSEC    bj. NS DS RRSIG NSEC
biz.                    86400   IN      RRSIG   NSEC 8 1 86400 20130724000000 201307
bj.                     172800  IN      NS      bj.cctld.authdns.ripe.net.
bj.                     172800  IN      NS      bow.rain.Fr.
bj.                     172800  IN      NS      ns1.intnet.bj.
bj.                     172800  IN      NS      ns-bj.afrinic.net.
bj.                     172800  IN      NS      nakayo.leland.bj.
ns1.intnet.bj.          172800  IN      A       81.91.225.18
nakayo.leland.bj.       172800  IN      A       81.91.225.1
bj.                     86400   IN      NSEC    bm. NS RRSIG NSEC
bj.                     86400   IN      RRSIG   NSEC 8 1 86400 20130724000000 201307
bm.                     172800  IN      NS      ns.uu.net.
bm.                     172800  IN      NS      ns1.bm.
```

Measurement Process

- Wait >10 seconds between queries per resolver
- If open resolver times out, retry (up to 5 times)

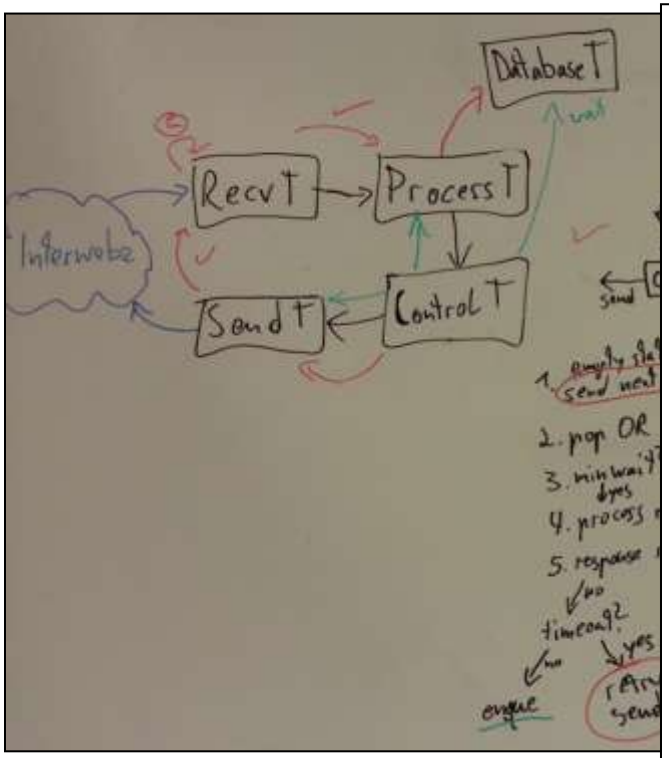

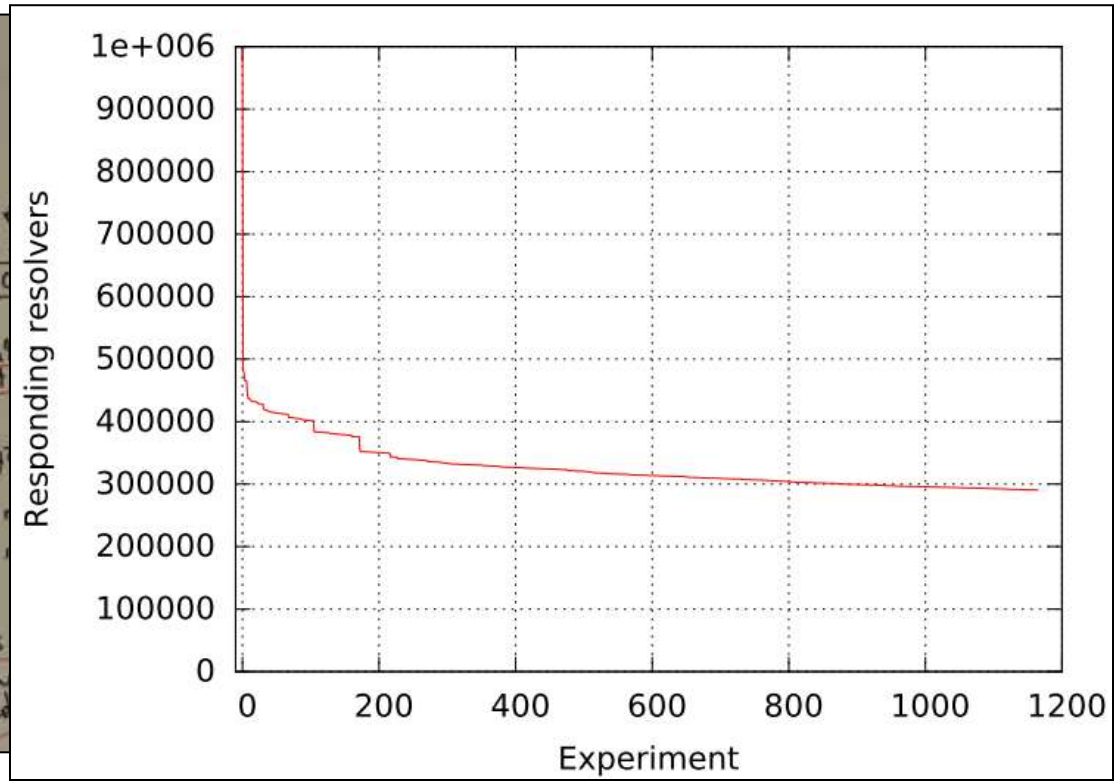

Results: Affected Open Resolvers

- 15k OR affected by Chinese DNS injection

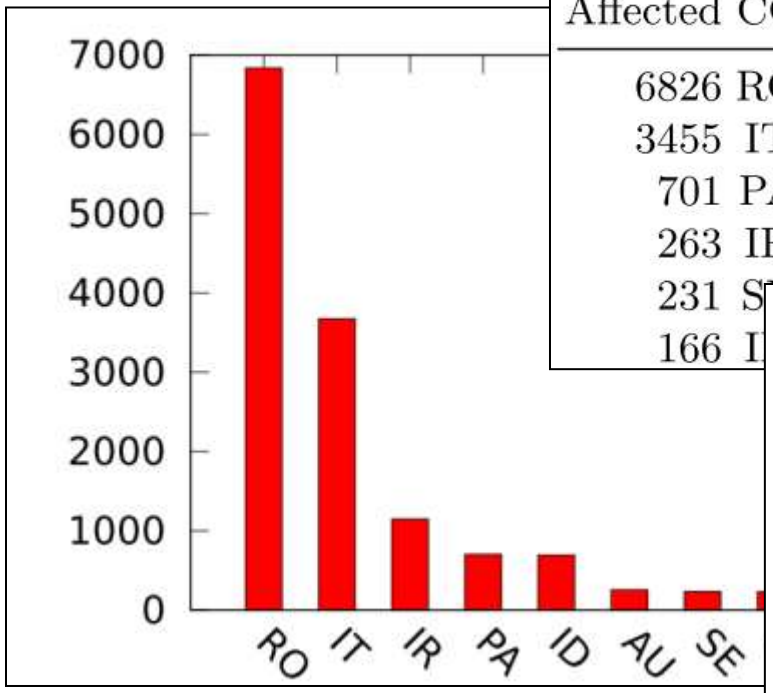

Affected	CC	AS#	Organization
6826	RO	9050	Romtelecom
3455	IT	3269	Telecom Italia
701	PA	11556	C&W Panama
263	IR	12880	ITC Iran
231	S		
166	I		

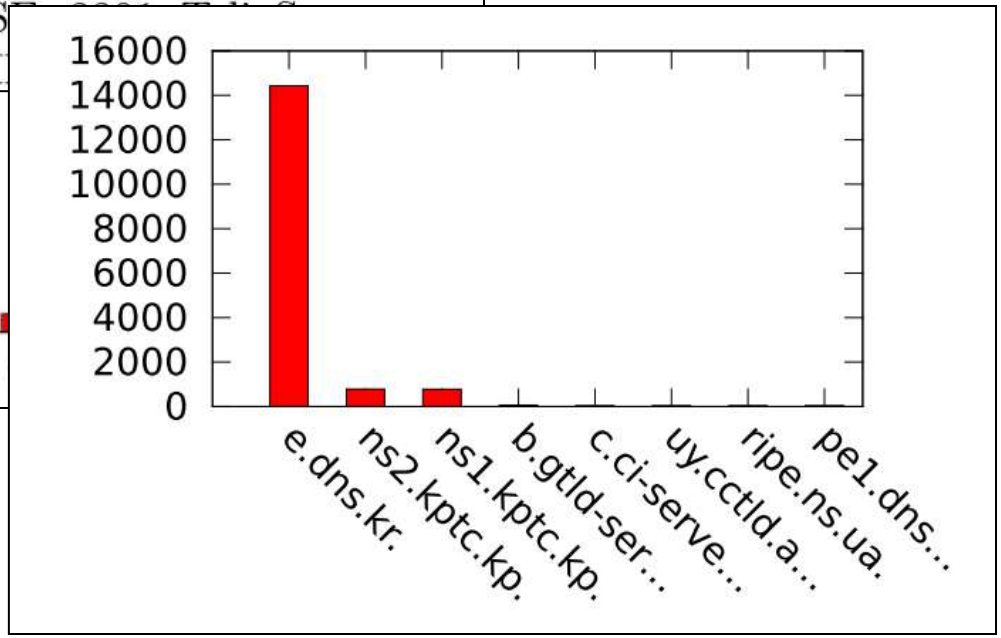

Complaints Received

- Probing IPv4 address space
  - Duisburg NOC: suspected malware
  - Kassel NOC: suspected UDP/53 portscan
  - None from destination networks (notified our NOCs?)
- Querying open resolvers
  - TLD operator 1: informed us about possible attack
  - TLD operator 2: forbid measurement with their server
  - None from operators of open resolvers

UNIVERSITÄT DUISBURG ESSEN

Analysis of Public DNSSEC Keys

Research Project 4

Analysis of Public DNSSEC Keys

- Objective: analysis of DNSSEC key material
  - Algorithms? Key lengths? (cf. SSL Observatory by EFF)
  - Easily factorable RSA keys? (cf. factorable.net)
- How to gather large amount of public keys?
  - Crawl DNSSEC zones by breaking NSEC3 hashes
  - "There is no name X with 3RL20V $< h(X) <$ 6D6ESU"

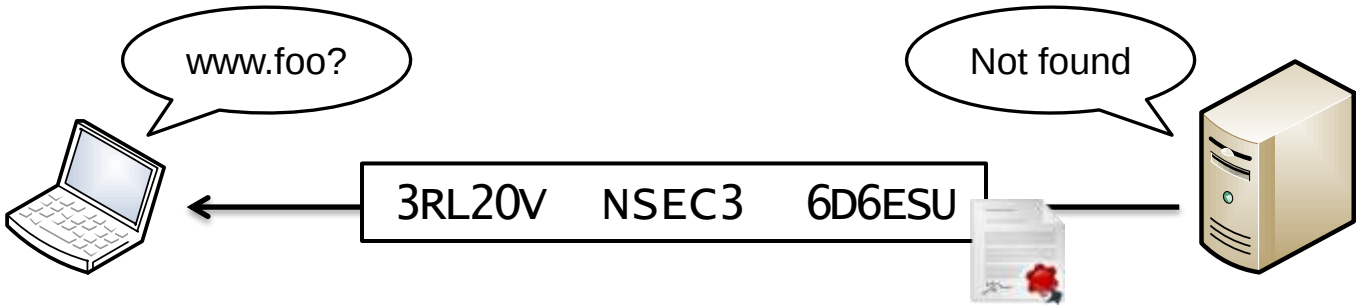

UNIVERSITÄT DUISBURG ESSEN

Effectiveness of DNS Caches

Research Project 5

Effectiveness of DNS Caches

- Network carrier provides DNS resolver + cache
  - Less latency for client name resolution?
  - Less load on authorative name server?

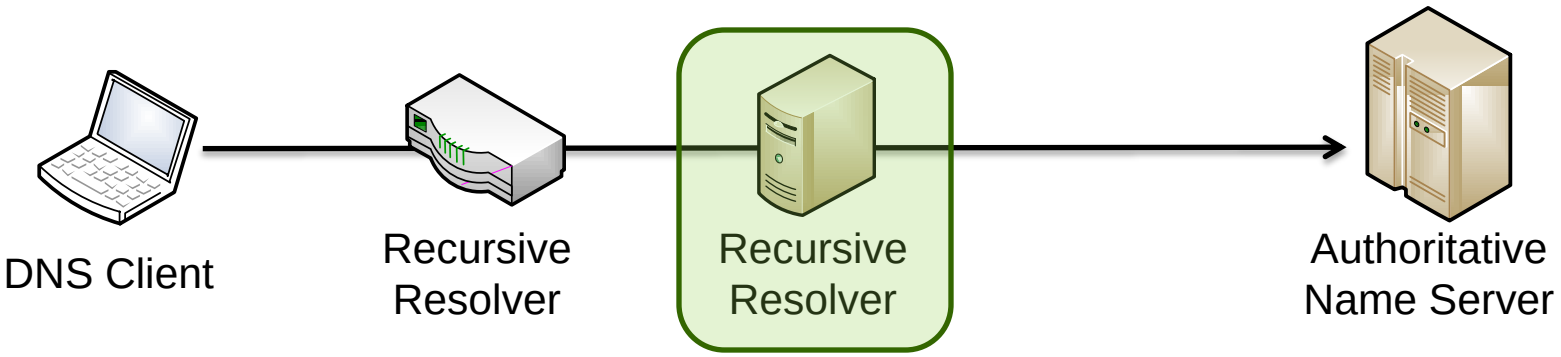

- How? Packet capture at NOC campus resolvers
  - Privacy! IP address ⇔ resolved domain names

Anonymization of Network Traces

- Script rewrites IP addresses from pcap file
  - `tcpdump -w - | python anonymize-pcap.py - anon-file.pcap`
- Set IP addresses to zero?
  - Need to distinguish clients
- Save hashed addresses?
  - Brute-force address space
- Keyed-hashing $h_K(\text{ip})$
  - Correlation attacks

```
<img src="http://uniq28384">
```

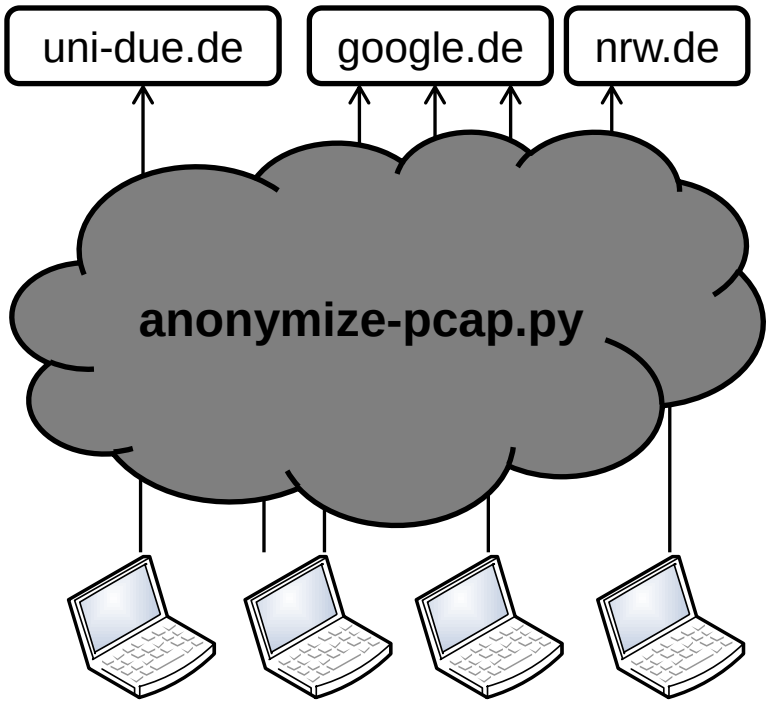

UNIVERSITÄT DUISBURG ESSEN

Analysis of Darknet Traffic

Research Project 6

Analysis of Darknet Traffic

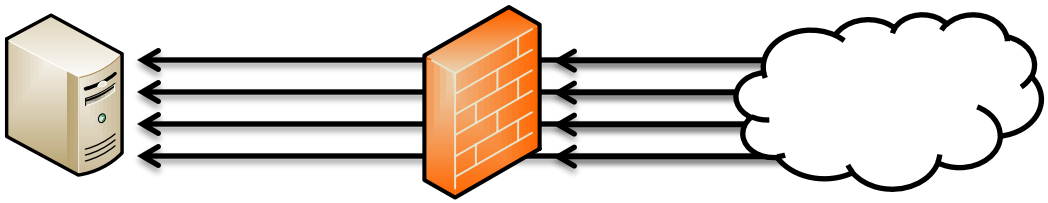

- Internet traffic arrives for unused IP addresses
- Objective: analysis of unsolicited Internet traffic
  - Assign addresses to host with packet capture
- Don‘t respond to incoming data, except for TCP:
  - Respond with SYN/ACK or ACK to get TCP payload
- Anonymization of unwanted traffic required?

Research Projects

1. Determine behavior of NAT routers — active
2. Count DNSSEC validating clients — active
3. Global impact of DNS censorship — active
4. Analysis of public DNSSEC keys — active
5. Effectiveness of DNS caches — passive
6. Analysis of darknet traffic — passive