Identifying Patterns in DNS Traffic

Pieter Lexis

System and Network Engineering

Thu, Jul 4 2013

Reflection and Amplification Attacks

- DNS abused as DDoS Tool
- Spamhaus hit with 300 Gigabit/second DDoS
- Reflected Amplification Attack
  - Send DNS query with spoofed source address to name server
  - Name server replies with a large(r) message to the victim
  - Flood the link to the victim

Reflection and Amplification Attacks

Prevention

- Firewalling on simple patterns
- BCP 38 (Network Ingress Filtering) [3]

Resolvers

- RFC 5358 ("Preventing Use of Recursive Nameservers in Reflector Attacks") [1]
- Firewalling based on IP addresses

Authoritative

- Response Rate Limiting [10]
  - Most Promising
  - Doesn't block all attacks [8]
- DNS Dampening [2]

Research Question

- How to analyse a large data set of DNS messages?
- How to recognize patterns in the data?

What types of behaviour can be detected in traffic to and from authoritative DNS servers and how can this detection be used to mitigate denial-of-service attacks?

Visualization

- Means of exploring data
- Uses the cognitive system to identify patterns
- Several visualizations for name server statistics exist
- Used before on resolver logs to identify security issues [6]

Data

- Packet Captures from authoritative name server from SURFnet
- 5 days of data
- 250 Gigabytes
- 630 million records
- Convert to JSON
- Inserted into ElasticSearch cluster


```
{
  "dns": {
    "additional": [],
    "answer": [],
    "authority": [],
    "edns": {
      "bufsize": 4096,
      "flags": {
        "DO": true
      },
      "version": 0
    },
    "flags": {
      "CD": true
    },
    "opcode": "QUERY",
    "qid": 34314,
    "question": [
      {
        "name": "ns1.surfnet.nl.",
        "type": "AAAA"
      }
    ],
    "rcode": "NOERROR"
  },
  "dport": 53,
  "dst": "192.87.106.101",
  "sport": 55564,
  "src": "203.0.113.77",
  "timestamp_unix": 1370304171.488599,
  "udp_len": 51
}
```

Tools

Rationale

Visual Information-Seeking Mantra

"Overview first, zoom and filter, then details-on-demand." [9]

- Batch tools
- Interactive GUI tools

Tools

Batch Tool 1 – Source Port versus Query ID

RFC 5452 (excerpt)[4]

```
 ...
Resolver implementations MUST:
 o Use an unpredictable source port for outgoing
queries from the range of available ports (53, or 1024
and above) that is as large as possible and
practicable;
...
 o Use an unpredictable query ID for outgoing queries,
utilizing the full range available (0-65535).
...
```

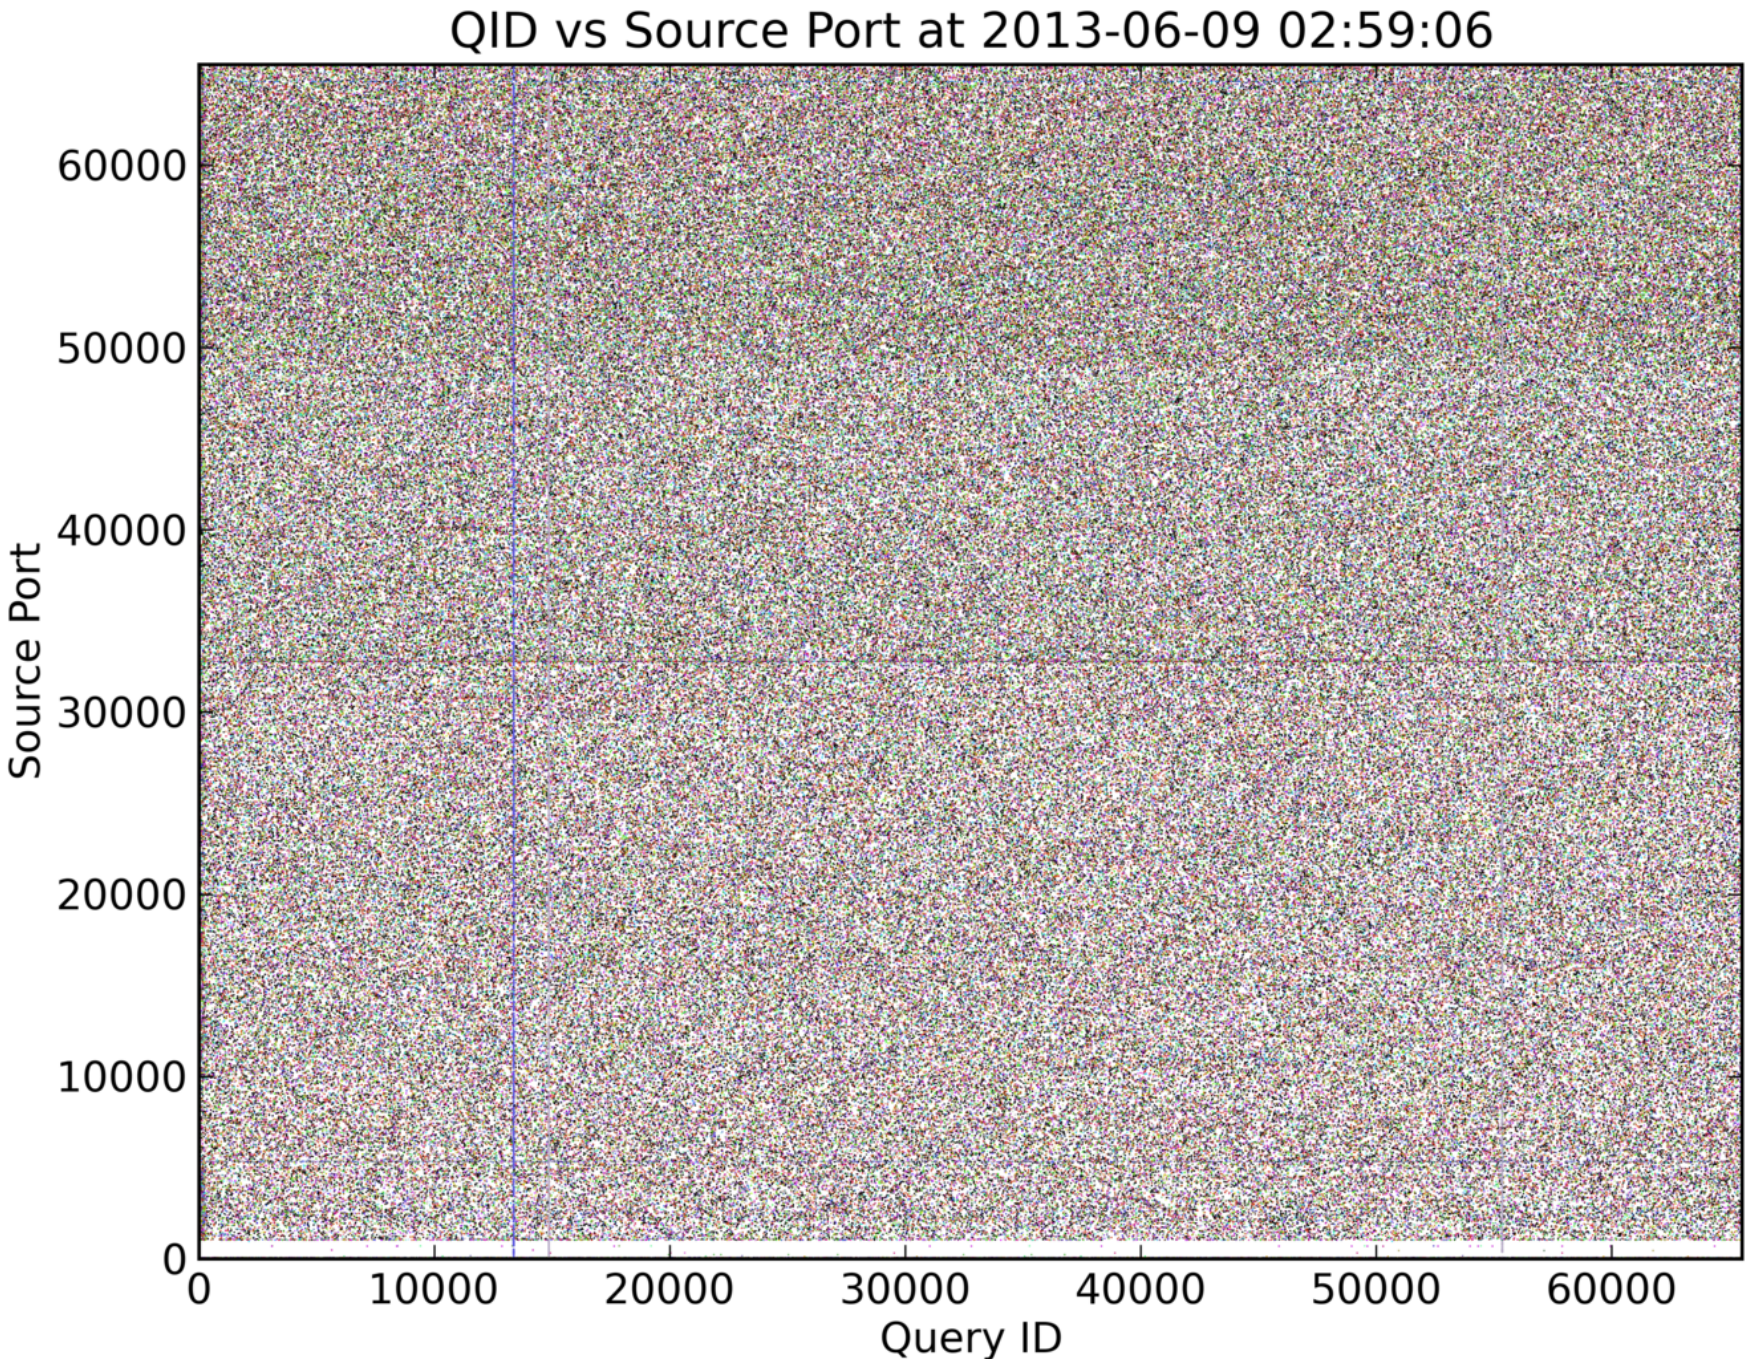
QID vs Source Port at 2013-06-09 02:59:06
Source Port
0
10000
20000
30000
40000
50000
60000
Query ID
0
10000
20000
30000
40000
50000
60000

Tools

Source Port Findings

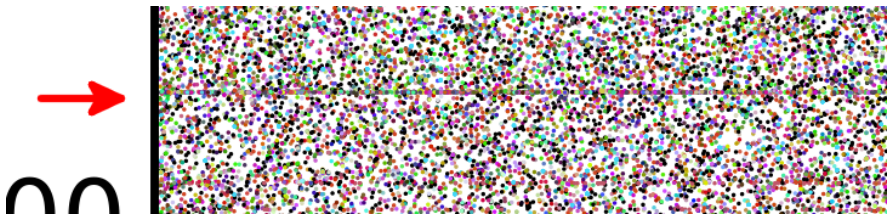

Bias of port numbers near 32768 (2^{15})

- Not a single source
- NAT Firewall

Increases ease of cache-poisoning attacks [5]

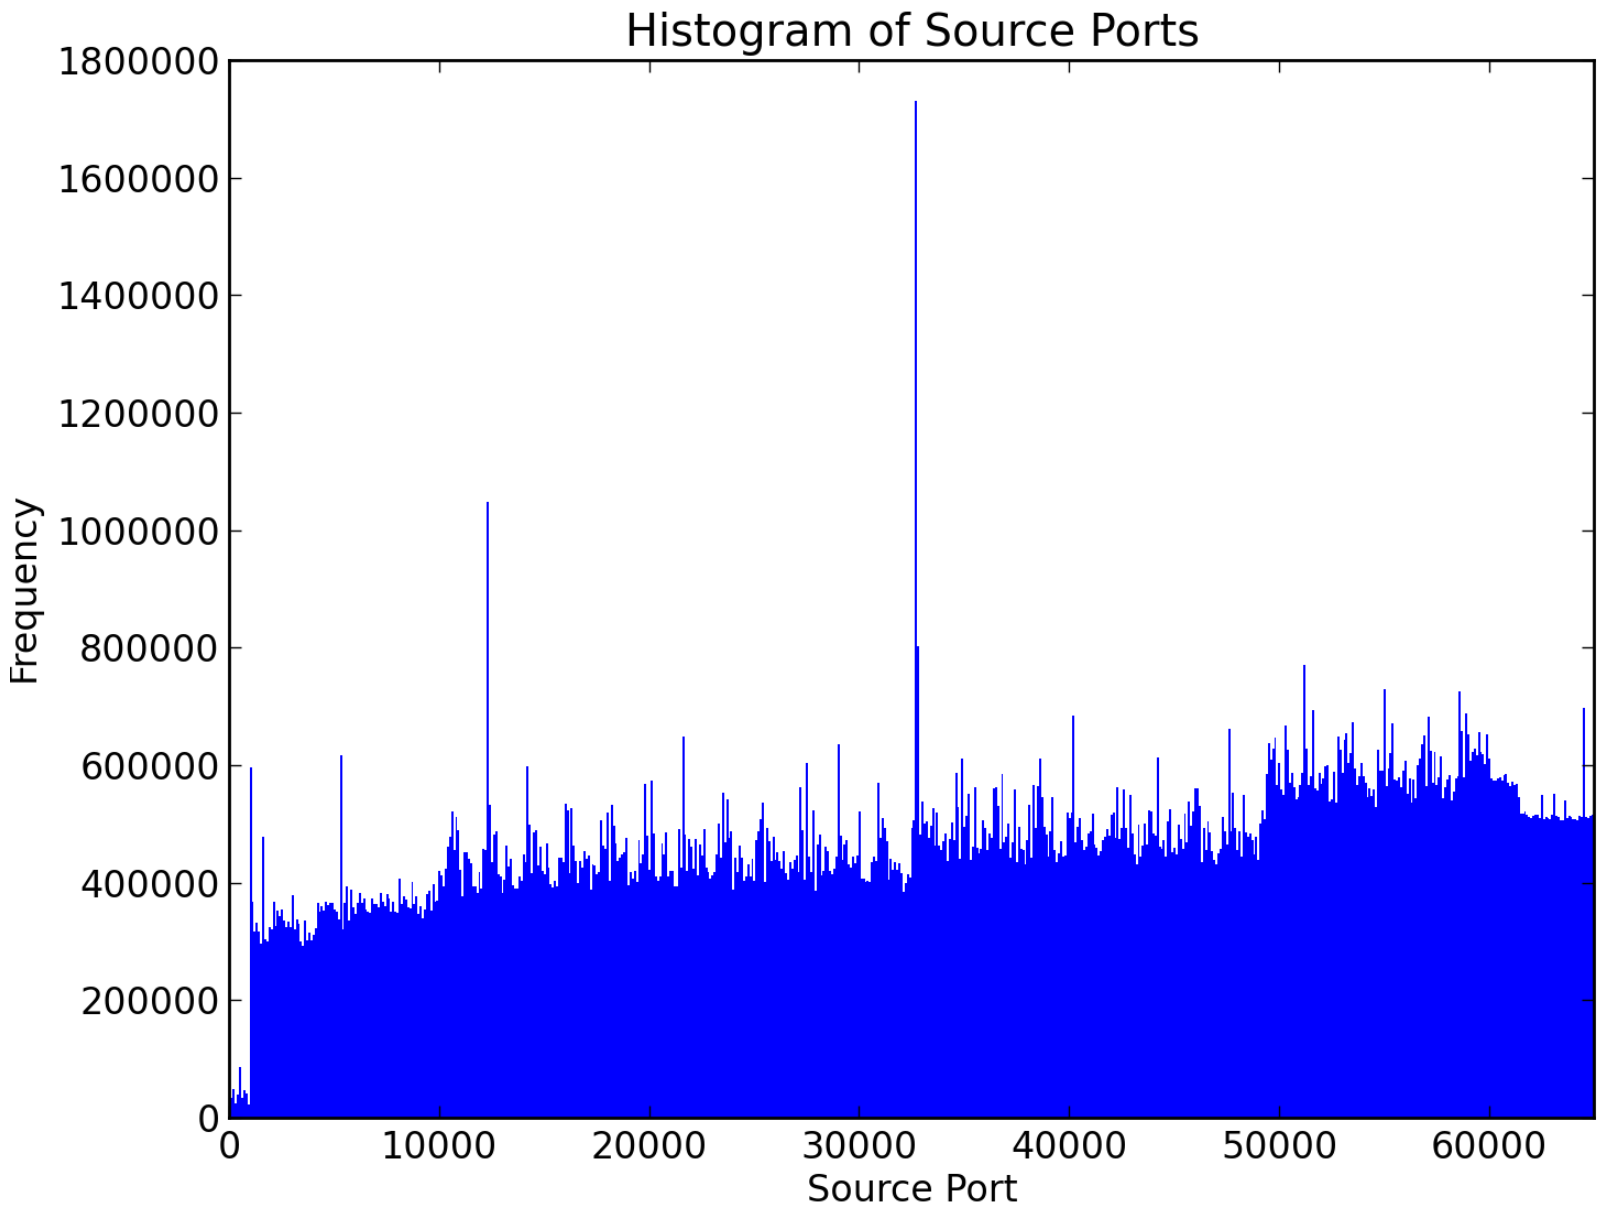
Histogram of Source Ports
Frequency
Source Port
0
200000
400000
600000
800000
1000000
1200000
1400000
1600000
1800000
0
10000
20000
30000
40000
50000
60000

Tools

Findings – Attack Spreading to Defeat Response Rate Limiting

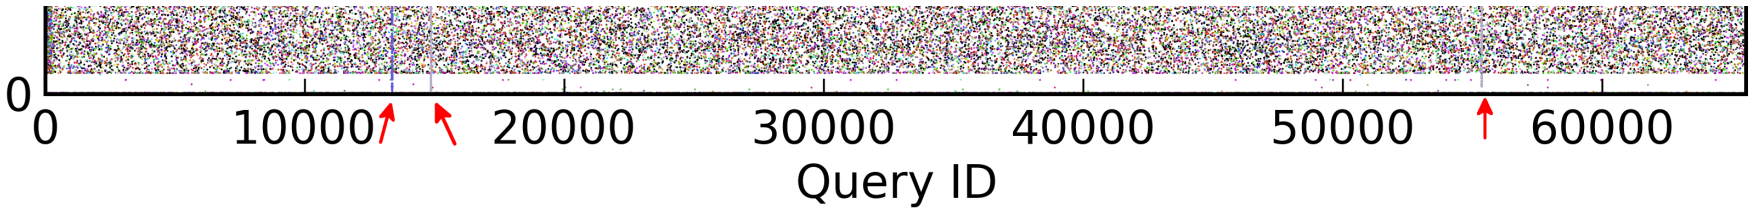

- Bias in Query IDs
- Queries are mostly `ANY`
- Query Names spread fairly evenly
- IP Addresses from a "DDOS protected" hoster

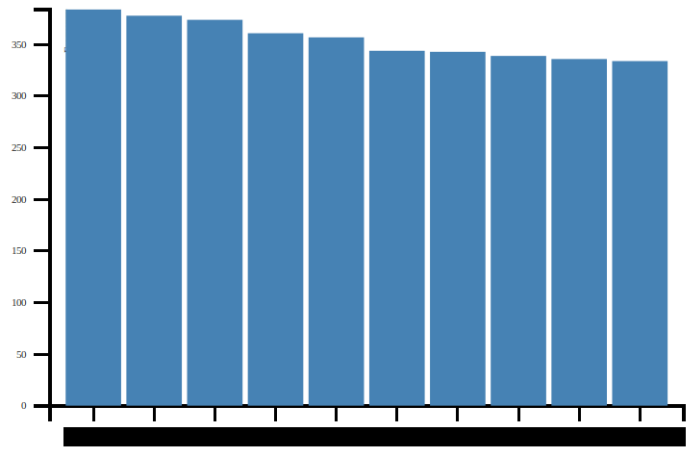

Figure: A bargraph with the frequency of Query Names for this IP Address.

Tools

Batch Tool – EDNS0 Buffersize Average

High buffersizes

- Might indicate abuse (large buffer → large response)
- Can cause fragmentation [7]

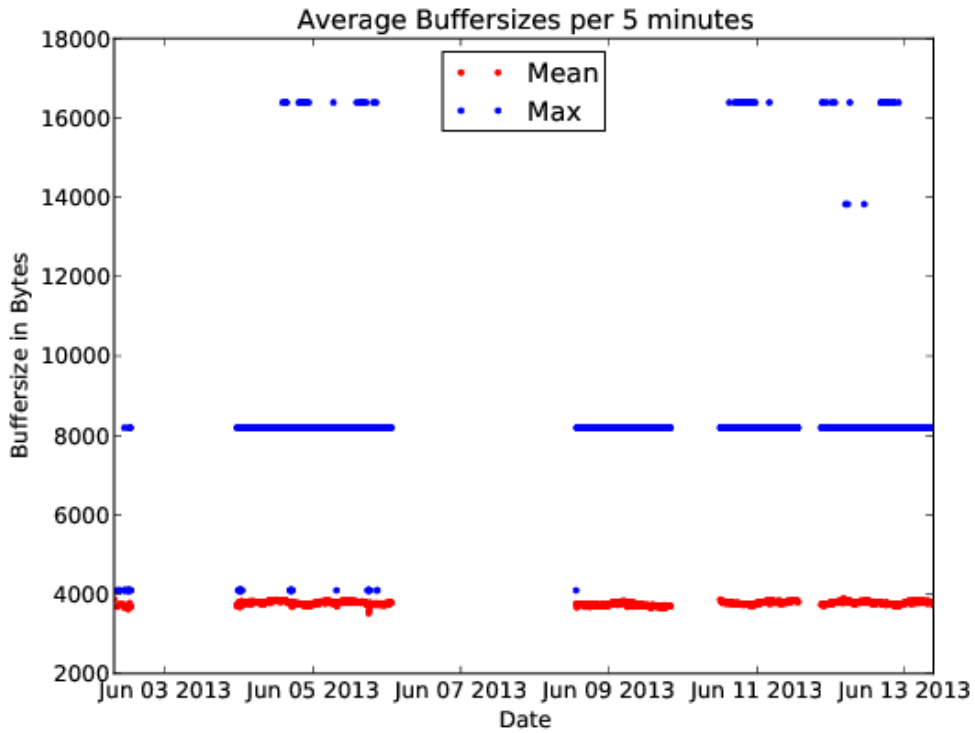

Tools

Interactive Tools

- Show data matching filters
- Filter on many of the fields/flags
- Used to zoom into the data

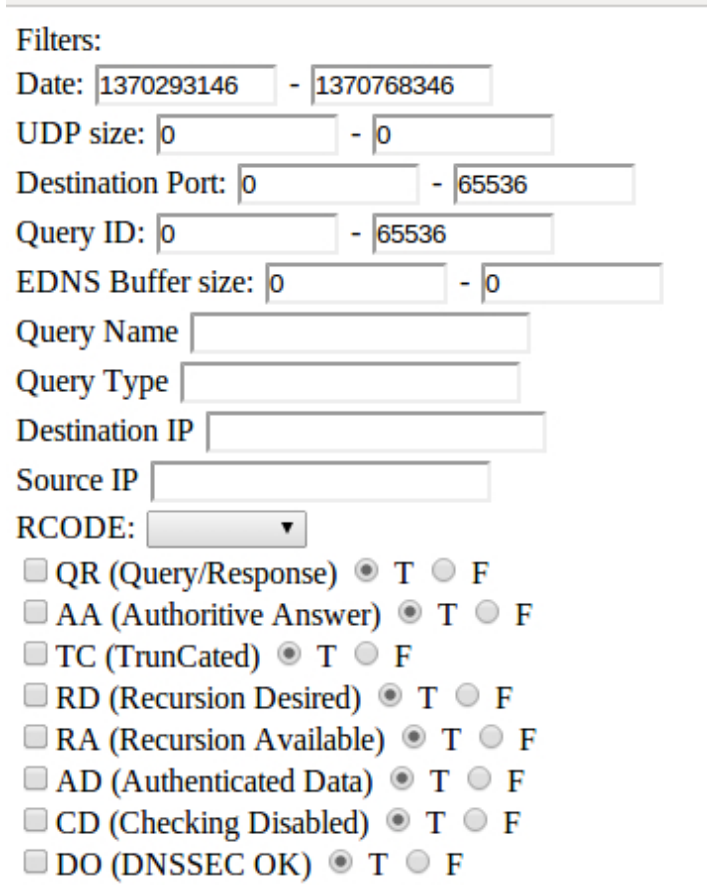

Tools

Interactive Tools – Aggregated View

- Frequency of values a field
- Keeps the previous graph + filters on-screen

Movie

Tools

Interactive Tools – Parallel Coordinates

- Shows the relationship between fields in messages
- Select fields to show
- Re-order axes
- Show subselections of axes

Movie

Conclusion

What types of anomalous behaviour can be detected in traffic to and from authoritative DNS servers and how can this detection be used to mitigate denial-of-service attacks?

- Several different anomalous behaviours detected
  - Source port selection of resolvers is not distributed well
  - Some attackers re-use query IDs
  - There are attacks in the wild that defeat RRL
- Visual approach works for initial identification, the insights gained could be used to develop new mitigation mechanisms

Future Work

- More interactivity
- Details on demand
- Real-time tools
- Statistical analysis of visually identified patterns
- Analyse more DNS message fields

QUESTIONS?

Bibliography (1)

[1] J. Damas and F. Neves. *Preventing Use of Recursive Nameservers in Reflector Attacks*. RFC 5358 (Best Current Practice). Internet Engineering Task Force, Oct. 2008. URL: http://www.ietf.org/rfc/rfc5358.txt.

[2] Lutz Donnerhacke. *DNS Dampening*. Accessed: 19 Jun 2013. 23 Sept 2012. URL: http://lutz.donnerhacke.de/eng/Blog/DNS-Dampening.

[3] P. Ferguson and D. Senie. *Network Ingress Filtering: Defeating Denial of Service Attacks which employ IP Source Address Spoofing*. RFC 2827 (Best Current Practice). Updated by RFC 3704. Internet Engineering Task Force, May 2000. URL: http://www.ietf.org/rfc/rfc2827.txt.

Bibliography (2)

[4] A. Hubert and R. van Mook. *Measures for Making DNS More Resilient against Forged Answers*. RFC 5452 (Proposed Standard). Internet Engineering Task Force, Jan. 2009. URL: `http://www.ietf.org/rfc/rfc5452.txt`.

[5] Dan Kaminsky. "Black ops 2008: It's the end of the cache as we know it". In: *Black Hat USA* (2008).

[6] Pin Ren, John Kristoff, and Bruce Gooch. "Visualizing DNS traffic". In: *Proceedings of the 3rd international workshop on Visualization for computer security*. ACM. 2006, pp. 23–30.

Bibliography (3)

[7] Roland van Rijswijk Deij. *DNSSEC and Fragmentation – A Prickly Combination*. Given at ICANN 45 in Toronto, 17 Oct 2012. 2012. URL: http://toronto45.icann.org/meetings/toronto2012/presentation-dnssec-fragmentation-17oct12-en.pdf.

[8] T Rozekrans and J de Koning. *Defending against DNS reflection amplification attacks*. 2013. URL: http://www.nlnetlabs.nl/downloads/publications/report-rrl-dekoning-rozekrans.pdf.

[9] Ben Shneiderman. "The eyes have it: A task by data type taxonomy for information visualizations". In: *Visual Languages, 1996. Proceedings., IEEE Symposium on*. IEEE. 1996, pp. 336–343.

Bibliography (4)

[10] Paul Vixie and Vernon Schryver. *DNS Response Rate Limiting (DNS RRL)*. URL: http://ss.vix.su/~vixie/isc-tn-2012-1.txt.