Measuring the DNS

Seminar Internet Measurements, TU-Berlin

Florian Streibelt

<florian.streibelt@TU-Berlin.DE>

July 20, 2012

Contents

Motivation

- Why is DNS so important?
- What is it good for?
- Who uses it?

Contents

Technically speaking...

- DNS maps hostnames to IP-addresses and vice versa
- until 1983: hosts file
- RFC 882 and RFC 883
- not only address mapping
- data organized in Resource-Records
- different types of RR exist: A, MX, TXT, NS, SOA, ...
- a DNS query consists of a chain of sections
- usually transmitted via UDP
- TCP is used for big queries, e.g. zone transfers
- characterization: globally distributed database

A simple DNS-Lookup

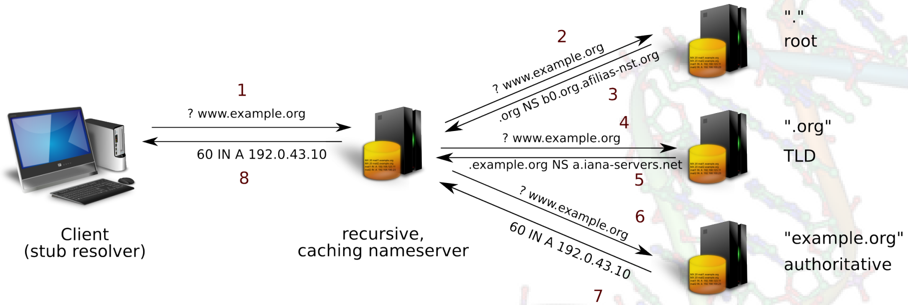

How DNS works

- client: stub resolver
- asks dedicated server
- often embedded in DSL Router
- usually forwards query to ISP
- builds a hierarchy of nameservers
- delegations exist between root, TLDs, domains, ...

(see next slide)

The DNS hierarchy

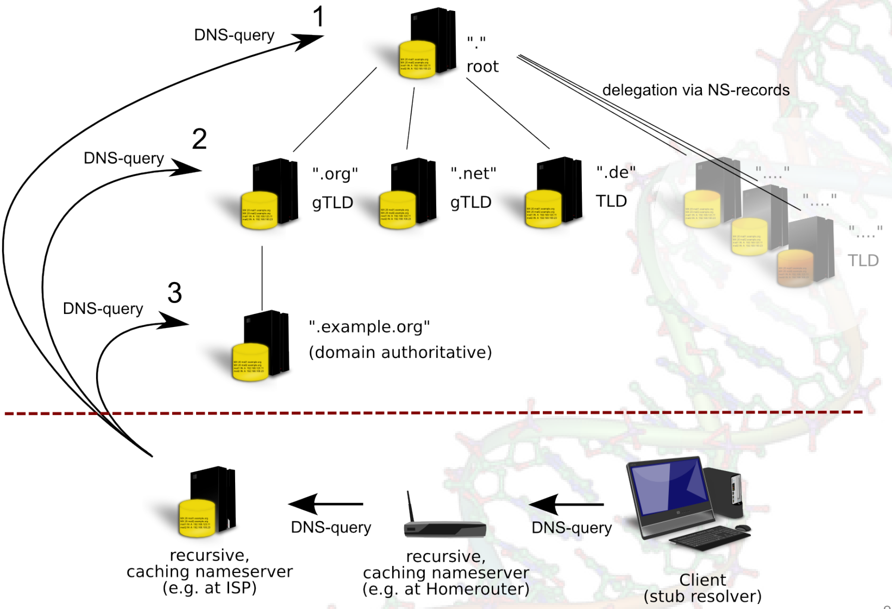

DNS-(Ab)uses today

- CDNs[1]: Records with low TTL
- used for traffic management, multiple levels of indirection
- ssh-key-fingerprints
- google site verification via TXT RRs
- Anti-Spam solutions like SPF, also TXT

$\Rightarrow$ partly considered an abuse, e.g. by P. Vixie [7]

[1] Content Distribution Networks

Contents

Considerations beforehand

- How can we measure?
- What can we measure?
- Where can we measure?
- Where should we measure?
- Why do we measure?

What to measure

- response time
- packet size
- number of queries
- cache hit rate and utilization
- valid/invalid requests
- success/failure rate
- types of queries
- ...

Where to measure

- application level (e.g. response time)
- client machine (e.g. number of failed/successful queries)
- customer router (e.g. cache utilization, response time, ...)
- provider (e.g. cache hit rate, failed/successful queries,...)
- root zone (e.g. illegal requests, wrong TLDs, ...)
- TLD (e.g. typo domains, number of clients,...)
- 2nd level domain (e.g. dos attacks, v4/v6 ratio, ...)

Levels in the DNS

- stub resolver at the client
- customer cache in the broadband router
- server at the ISP or company
- domain server
- TLD servers
- root servers
- New in the game: dedicated public DNS providers (Google, OpenDNS)

⇒ Each level is special in the data observable!

Contents

Client Level (I)

Considerations:

- at the bottom level
- maximum 'propagation delay':
- requests have to travel trough whole DNS
- all layers are caching
- impact of these caches?

Client Level (II)

"DNS Performance and the Effectiveness of Caching" [5]:

- traces collected at two large universities
- 23% of the queries returned no answer
  - responsible for more than half of the DNS-packets:
  - retransmission of failed requests by clients!
- most frequently requested query types:
  - A records (60%), Hostnames to Addresses
  - PTR, reverse lookups (24% to 31%)

Client Level (III)

"DNS Performance and the Effectiveness of Caching" [5]:

- latency of queries
  - number of referrals: big impact
  - for every referral another nameserver has to be contacted
- 70% returned direct answers from cache
- impact of caching can be observed (latency)
- between 10% and 42% result in a negative answer,
- explanation: no reverse-mapping, server misconfiguration, other failures

Client Level (IV)

"DNS Performance and the Effectiveness of Caching" [5]:

- idea: negative caching could have impact on overall performance
- but: heavy-tailed distribution of names
- the effect of negative caching is limited

- idea: negative caching could have impact on overall performance
- again: heavy-tailed nature of access to names!
  - but TTL for NS-records is reasonable high (cached)
  - use of dynamic low-TTL A-record bindings should not degrade DNS performance
- disagrees with Vixie in [7]

ISP-Level

"An analysis of wide-area name server traffic[...]" [4]:
Considerations:

- aggregates queries from several users
- currently their location is used for geolocation
- traditionally: fast servers with huge caches

Observations:

- cache hit rate was over estimated (last slides)
- but can serve as 'firewall' between WAN/root DNS and malicious or badly broken programs/users

TLD-Level (I)

- 1 of 13 nameservers for `.com` and `.net` (`g.gtld-servers.net`)
- measurements below the rootservers
- passive measurements (packet dumps)

Expectations by Osterweil et al. [6]:

- audience for the generic TLDs is not country-specific
- resolvers will probe each authoritative server of a zone over time
- nearly all resolvers in the Internet should be seen
- should also see more variance than a rootserver

TLD-Level (II)

Observations:

- traffic from the whole active IP-Range of the Internet is being observed,
- most prominently asked: `A` (70%)[2], `AAAA` (15%) and `MX` (10%).
- also deprecated types are still being queried, e.g. `A6` records, superseeded by `AAAA`

[2] all approximately values derived from a figure

TLD-Level (III): Top Talkers

Main Aspect of their work:

- small set of 40.000 resolvers, called top-talkers, is responsible for 90% of overall traffic
- the set of top-talkers is highly evolving
- the same set of clients only accounts for 84% percent of the queries after 10 days,
- a huge number of new clients is constantly detected.
- list of top-talkers can be used as a filter in large-scale measurements to dramatically reduce the amount of datapoints, with small effect on accuracy of the measurements.
- yet unclear if a global set of top-talkers can be identified

Informal summary: “Heavy tail is everywhere.”

Root-Level (I)

Considerations:

- up-most level in the hierarchy of nameservers.
- these serves act as last resort if no authority can be found for a subtree
- hit by a lot of attacks (actively and by misconfigurations)
- key element of the Internet (see latest discussions on control over them)
- without them as single root: no anchor, no determinism
- when they fail the Internet will break for end users after caches time out

Root-Level (II)

"DNS measurements at a root server" [2]

- massive amount of erroneous queries
- not always clear, if queries are result of a misconfiguration or attack
- "60-85% of observed queries were repeated from the same host within the measurement interval."
- "Over 14% of a root server's query load is due to queries that violate the DNS specification."
- abused for amplifiaction/reflection attacks using forged source addresses (UDP)

Root-Level (III)

Castro et. al. in [3]

- "an estimated 98% of the traffic at the root servers should not be there at all"
- very high level of DNS pollution.

main findings:

- most of the traffic could be avoided by proper **caching**
- huge amount of traffic originates at **improperly configured** resolvers
- outstanding example:
  - one client querying a non-existent SOA record over 2 million times in an hour
  - caching would have reduced this to less than 10.

Contents

EDNS

- pseudo-RR was needed, because header had no bits left
- packets getting bigger
- sometimes to big for UDP
- usage of EDNS is increasing:
  - DNSSEC (cryptographically signed DNS data)
  - client IP transmission (for geolocation)

DNSSEC - What and Why

- provides integrity for all replies
- Why? SSH-key fingerprints, e-mail transportation, Online Banking, Domain stealing (ebay)
- each record has to be signed individually, as each one is individually requested.
- cryptographic algorithms are used must be compatible with caching.

DNSSEC - Impacts?

Considerations:

- operation will become more complex, making human error more likely
- bandwidth used will increase because signatures and keys have to be transmitted
- caches will have to hold more information
- validating resolvers will have to do complex cryptographic operations to validate each response.

DNSSEC - Measurements

"Exploring the overhead of dnssec" [1]:

- simulation based on measurement data and collected traces
- for RSA-signatures **packet size** grows "on average by a factor of 3.4 and 12.7 in the worst case."
- can lead to a fallback to TCP for the query, adding 5 rtts
- **memory requirements** of nameservers are significantly higher, factor of four for caching NS

Also:

- home-routers: overflow of resources was observed[3]
- impact on e.g. root servers is not yet totally clear

[3] `https://lists.dns-oarc.net/pipermail/dns-operations/2010-September/006123.html`

Contents

Summary - Caching

Effects of caching had been incorrectly estimated:

- main effect is not caching of positive A-answers (heavy tail!)
- help to avoid unnecessary lookups to find authoritative servers
- help by using negative caching to prevent repeated queries
- small TTLs as used by CDNs have no big impact
- positive side-effect of caching: robustness against attacks and misconfigured systems.

Summary - Stability

- DNS is more stable than expected.
- large amount of bogus traffic is not affecting the DNS as a whole due to working delegation and caching
- But: single incidents like the outage of RIPEs reverse DNS or Microsofts nameservers show, how a configuration error on one site still can affect the system as a whole.

Summary - Scalability

- more and more services depend on DNS and use it as a general purpose data-store or lookup service
- DNS as global address mapping service is working stable and still scaling well, partitioning into independently administered zones works

Summary - Issues

- amount of bogus traffic on the rootservers is still surprisingly high and should be continued to act upon.
- The exact effect of extensions like DNSSEC on the load of caching servers and the rootservers is unknown.
- worst case: overloading the servers, denial of service
- impact of initiatives like “A Faster Internet” can currently not be estimated, as there is no data publicly available.

Bibliography I

- Bernhard Ager, Holger Dreger, and Anja Feldmann.
  Exploring the overhead of dnssec.
  2005.
- N. Brownlee, K.C. Claffy, and E. Nemeth.
  Dns measurements at a root server.
  In *Global Telecommunications Conference, 2001. GLOBECOM '01. IEEE*, volume 3, pages 1672 –1676 vol.3, 2001.
- Sebastian Castro, Duane Wessels, Marina Fomenkov, and Kimberly Claffy.
  A day at the root of the internet.
  SIGCOMM Comput. Commun. Rev., 38(5):41–46, September 2008.

Bibliography II

- Peter B. Danzig, Katia Obraczka, and Anant Kumar.
  An analysis of wide-area name server traffic: a study of the internet domain name system.
  SIGCOMM Comput. Commun. Rev., 22(4):281–292, October 1992.
- Jaeyeon Jung, E. Sit, H. Balakrishnan, and R. Morris.
  Dns performance and the effectiveness of caching.
  Networking, IEEE/ACM Transactions on, 10(5):589 – 603, oct 2002.

Bibliography III

Eric Osterweil, Danny McPherson, Steve DiBenedetto, Christos Papadopoulos, and Dan Massey.
Behavior of dns' top talkers, a .com/.net view.
In *Lecture Notes in Computer Science 7192, Proceedings of 13th Internation Conference Passive and Active Measurement, PAM 2012*, PAM 13, pages 211–221, Berlin a.o., 2012. Springer.

Paul Vixie.
What dns is not.
Commun. ACM, 52(12):43–47, December 2009.

Download:
http://f-streibelt.de/talks/

Acknowledgments:
C.Fuerst for his reliable (email) reminders...
Enric Pujol for proof-reading the slides and paper

Image sources:
DNA: http://en.wikipedia.org/wiki/File:DNA_Overview2.png
others: own work and http://openclipart.org/