DNS & Iodine

Christian Grothoff
christian@grothoff.org
http://grothoff.org/christian/

"The Domain Name System is the Achilles heel of the Web."
– Tim Berners-Lee

DNS: Domain Name System

- Unique Distributed Database
- Application-layer protocol over UDP or TCP
- Maps names to IP addresses
- IP addresses to names
- Load distribution (multiple IP addresses for one canonical name)

Why not centralize DNS?

- single point of failure
- traffic volume
- high latency for those further away
- maintenance

$\Rightarrow$ Centralized does not scale!

Key DNS Services

- Hostname to IP address translation (A, AAAA)
- Host aliasing (canonical name, CNAME record)
- Mail server aliasing (MX records)
- Nameserver delegation (NS records)
- Arbitrary text (TXT records)

Distributed, Hierarchical Database

NS-records are used to specify delegations.

Root name servers[1]

[1] http://www.root-servers.org/map/

Top-Level Domain (TLD) Servers

- Responsible for com, org, net, edu, etc and top-level country-code domains (de, uk, fr, ca, jp, eu)
- Organizations hosting TLD servers:
  - de: DENIC
  - edu: Educause
  - com: Network Solutions
- These organizations perform “domain-name registration”

Authoritative DNS Servers

- Individual organization's DNS servers, providing **authoritative** mappings for organization's servers
- Can be maintained by organization or service provider
- Subdomains (www) and services (MX) are specified here
- Further delegation possible: `news.bbc.co.uk`

Local Name Server

- Does not strictly belong to hierarchy
- Each ISP (residental ISP, company, university) has at least one, typically at least two
- Also called “default name server”
- Hosts query the local DNS server, acts as proxy, forwards query into hierarchy

Stub Resolver

- DNS Resolver running on each host
- Often build deep into the OS
- Not a full DNS implementation
- Translates calls to `gethostbyname` or `getaddrinfo` into interaction with local name server

Iterative DNS Lookup

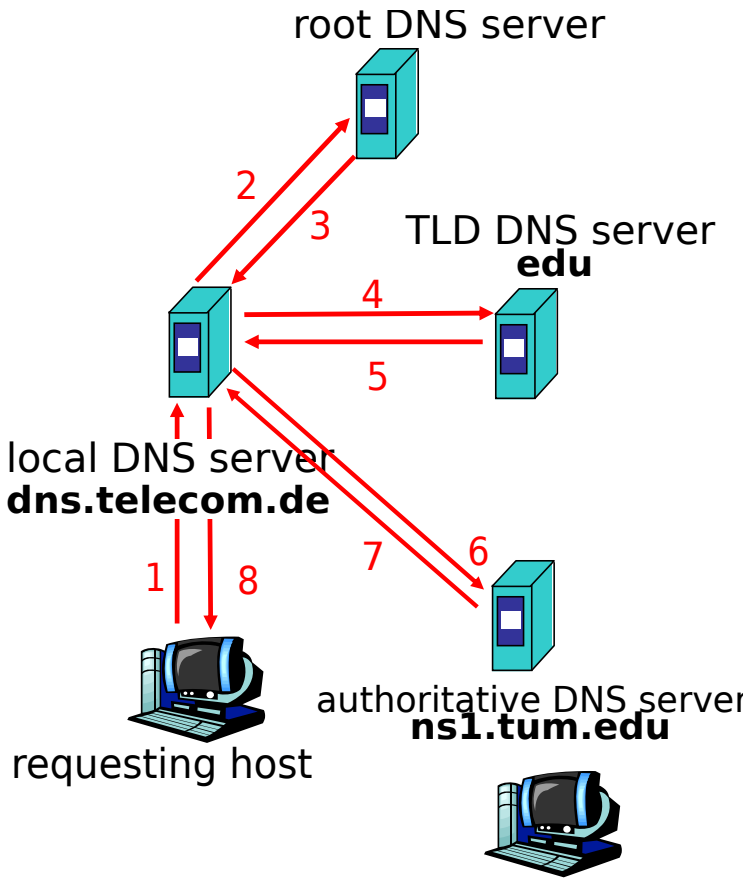

Recursive DNS Lookup

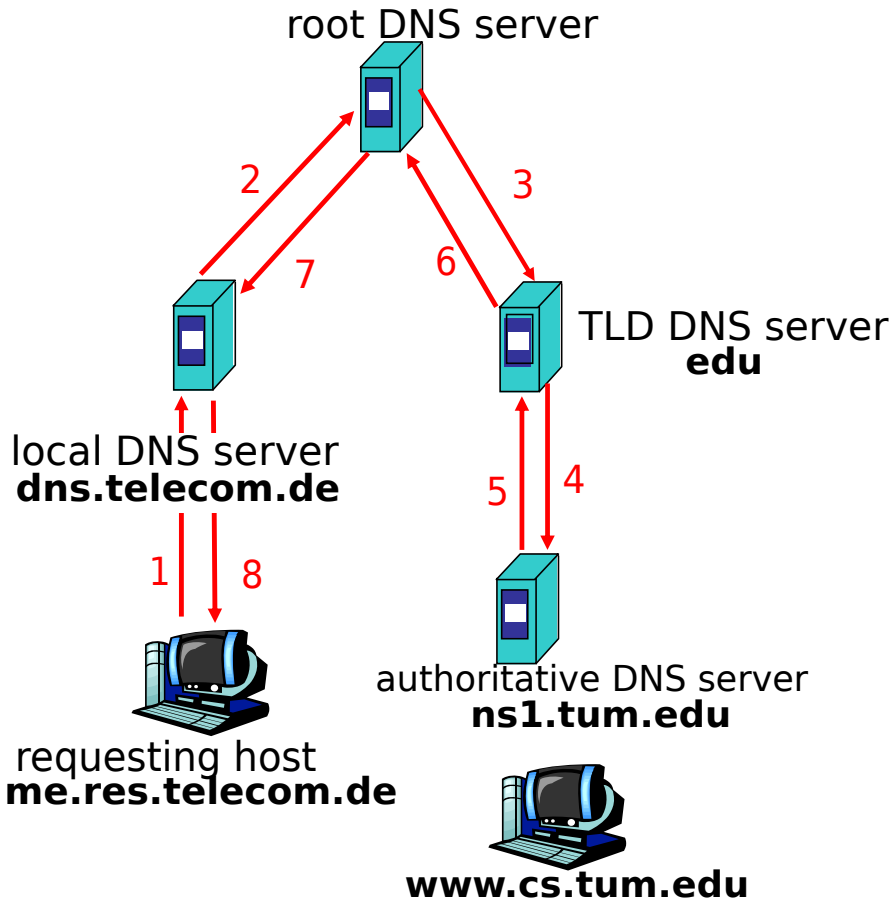

DNS Caching

- DNS servers cache mappings (DNS records)
- Entries time out based on expiration time in DNS record
- TLD servers are typically cached in local name servers

$\Rightarrow$ Root name servers not visited often

DNS Zone Transfers

- A DNS Zone transfer copies the DNS database
- Organization typically has “backup” (secodary) DNS server
- Changes to primary DNS database must be propagated to backup server
- RFC 2136 “DNS UPDATE” specifies incremental update for fast convergence

DNS Records

Records always contain four values:

- Name (a string)
- Value
- Type (a short string)
- Time-to-live (TTL) — how long caching is allowed

Each record originates from the authority for the respective name.

A records

- Name: hostname
- Value: IPv4 address
- Type: “A”

AAAA records

- Name: hostname
- Value: IPv6 address
- Type: “AAAA”

NS records

- Name: domain (e.g. foo.com)
- Value: hostname of authoritative name server
- Type: “NS”

MX records

- Name: domain (e.g. foo.com)
- Value: hostname of mail (SMTP) server
- Type: “MX”

CNAME records

- Name: alias name (i.e. www.ibm.com)
- Value: canonical (real) name (i.e. www2.eastcoast.ibm.com
- Type: “CNAME”

DNS Protocol (1/2)

Query and reply messages have the same format:

Identification	Flags
number of questions	number of answer RRs
number of authority RRs	number of additional RRs
questions (variable number)	
answers (variable number)	
authority (variable number)	
additional information (variable number)	

DNS Protocol (2/2)

- Identification: 16 bit number of query that reply must match
- Flags: query or reply, recursion desired / recursion available, reply is authoritative
- Questions: name, type of a query
- Answers: RRs in response to query
- Authority: records for authoritative servers (NS records)
- Additional information: RRs that might be useful as well

Inserting records into DNS

- Register "grothoff.org" at DNS registrar:
  - Provide name and IP of authoritative DNS server
  - Registrar inserts two RRs into org TLD server:
    $(grothoff.org, ns1.grothoff.org, NS)$
    $(ns1.grothoff.org, 12.34.56.78, A)$

- Configure authoriative server:
  $(www.grothoff.org, 12.34.56.79, A)$
  $(home.grothoff.org, 12.34.56.80, A)$
  $(home.grothoff.org, 2001:24::1, AAAA)$
  $(mail.grothoff.org, home.grothoff.org, CNAME)$
  $(grothoff.org, mail.grothoff.org, MX)$

DNS and IP Anycast

IP anycast makes multiple servers reachable under the same IP address:

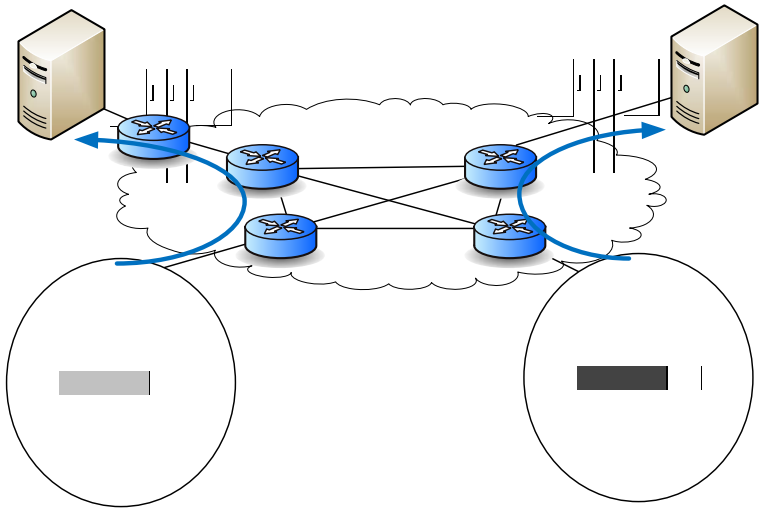

IP anycast is used for root servers and many TLDs since 2002.

fsnsg

Dependency on DNS

DoS-Attack targeted Microsoft in January 2011:

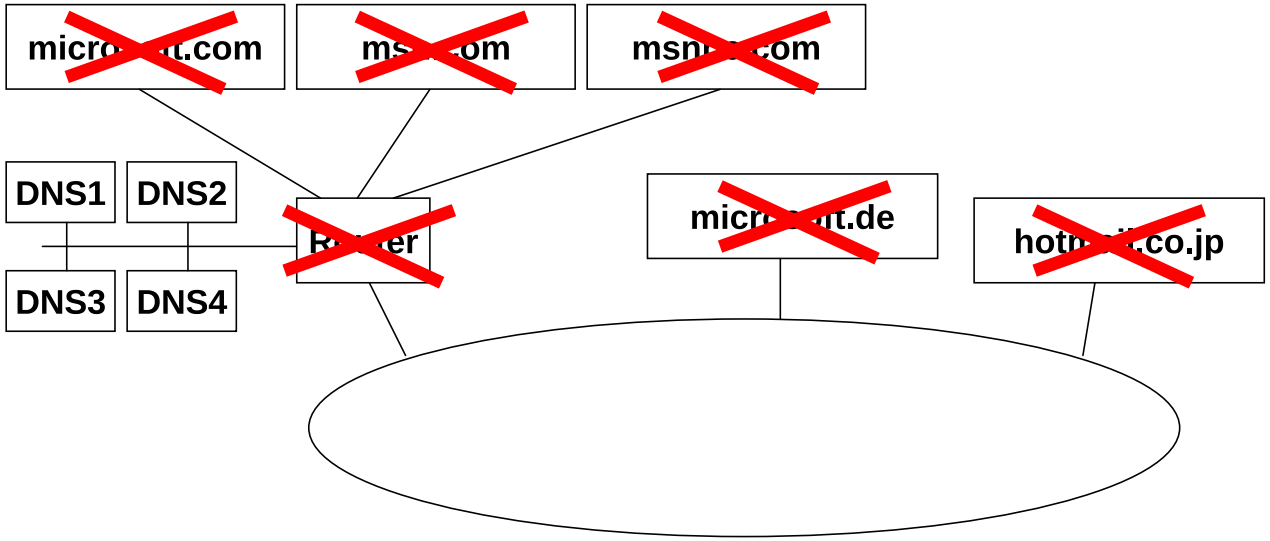

Questions

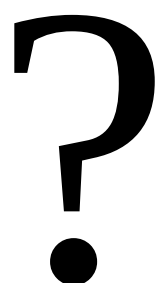

"The only truly secure system is one that is powered off, cast in a block of concrete and sealed in a lead-lined room with armed guards — and even then I have my doubts." –Gene Spafford

IP over DNS with Iodine

Christian Grothoff
christian@grothoff.org

"Never let your sense of morals get in the way of doing what's right." –Isaac Asimov

Disclaimer

This is an educational presentation:

- Ask your geek (me) how to do this
- Ask your lawyer about the legality of this
- Ask your priest about the ethics of this

Problem Statement

Your "provider" offers an open WLAN network with browser-based authentication and/or payment.

Specifically:

- Your local ISP gives you DNS, but not IP service
- "nslookup www.google.com" works prior to payment

How can you go online anyway?

Related Problem Statements

- You are at a university and the conference-provided username/password doesn't work, or
- The university is a bit insane and filters `ssh`, `irc` or other useful protocols

Solution

Tunnel IP over DNS

After all, DNS is allowed, right?

Prerequisites

- Iodine (client and server), seems portable
- Control (root) over some system
- Control over a domain (i.e. `grothoff.org`)

Setup (while at home...)

- Point “ns” record of “i.grothoff.org” to your machine (i.e. “my.home.in.tum.de”)
- Start “iodined” on “my.home.in.tum.de”:

```
# iodined -c -f -D -u grothoff \
  -P password 192.168.0.1 i.grothoff.org
```

Setup (while at home...)

- Configure NAT (on “my.home.in.tum.de”):

```
# echo 1 > /proc/sys/net/ipv4/ip_forward
# iptables -t nat -A POSTROUTING -o eth0 -j MASQUERADE
# iptables -A FORWARD -i eth0 -o dns0 -m state \
  --state RELATED,ESTABLISHED -j ACCEPT
# iptables -A FORWARD -i dns0 -o eth0 -j ACCEPT
```

Setup WLAN (on the road)

- For example:

```
# wpa_supplicant ...
# dhclient wlan0
```

- Check your ISP's DNS resolver:

```
# cat /etc/resolv.conf
nameserver 62.101.93.101
```

- Check your default route:

```
# route -n | grep 0.0.0.0
0.0.0.0  10.10.10.1 ...
```

Drop ISP's default route

- Remove your existing default route:

```
# route del default gw
```

- Keep routing DNS queries via old default route:

```
# route add -host 62.101.93.101 gw 10.10.10.1
```

Setup Iodine

- Start iodine:

```
# iodine -f -u grothoff -P password \
  -L0 62.101.93.101 i.grothoff.org
```

- Add new default route:

```
# route add default gw 192.168.0.1
```

A few tests: ping

```
# ping -c1  www.net.in.tum.de
PING www.net.in.tum.de (131.159.15.49) 56(84) bytes of dat
64 bytes from typo3.net.in.tum.de (131.159.15.49):
  icmp_req=1 ttl=62 time=27.6 ms

--- www.net.in.tum.de ping statistics ---
1 packets transmitted, 1 received, 0% packet loss, time 0m
rtt min/avg/max/mdev = 27.634/27.634/27.634/0.000 ms
```

A few tests: ssh

```
$ time ssh gnunet.org -C echo test
test

real 0m2.186s
user 0m0.048s
sys 0m0.040s
```

A few tests: wget

```
$ time wget -o /dev/null \
  http://grothoff.org/christian/

real 0m6.456s
user 0m0.000s
sys 0m0.020s
```

A few tests: scp

```
$ ls -al test.pdf
-rw-r--r-- 1 g g 303297 test.pdf

$ time scp test.pdf my.home.in.tum.de:.
real 1m0.900s
user 0m0.104s
sys 0m0.016s
```

So expect roughly 5 kb/s upload, I got about 15 kb/s downloads.

Experiences

If you only do one thing at a time, these work:

+ IMAPS
+ SMTP
+ HTTP / HTTPS
+ SSH

What's actually going on?

```
Source: 62.101.93.101 (62.101.93.101)
Destination: 10.10.10.33 (10.10.10.33)
User Datagram Protocol
Source port: domain (53)
Destination port: 38275 (38275)
Length: 752
Transaction ID: 0x12e5
Queries
    paaiglzq.i.grothoff.org: type NULL, class IN
    Name: paaiglzq.i.grothoff.org
    Type: NULL (Null resource record)
    Class: IN (0x0001)
```

What's actually going on?

```
Answers
 paaiglzq.i.grothoff.org: type NULL, class IN
 Name: paaiglzq.i.grothoff.org
 Type: NULL (Null resource record)
 Time to live: 0 seconds
 Data length: 631
 Data
Authoritative nameservers
 i.grothoff.org: type NS, class IN, ns my.home.in.tum.de
 Name: i.grothoff.org
 Type: NS (Authoritative name server)
 Name Server: my.home.in.tum.de
```

Questions

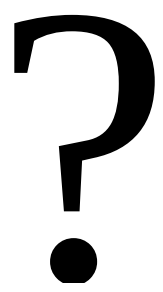

"The most all penetrating spirit before which will open the possibility of tilting not tables, but planets, is the spirit of free human inquiry. Believe only in that." – Dmitri Mendeleev

RTFL