

# Naming


Martin Heusse

# DNS

## Domain Name Service

- Initially designed for name to address translation
- `/etc/hosts` file : limited solution...(does not scale)  
192.168.1.10      foo.mydomain.org    foo  
Used to be manually distributed on hosts
- `/etc/resolv.conf` file (DNS servers) (since 1984)  
search imag.fr  
nameserver 129.88.38.1  
nameserver 129.88.30.10
- Name hierarchy, matching a server hierarchy
- DNS : spine of the web, also of e-mail...

# Naming hierarchy



- “fr” is a *Top-Level Domain*
  - A zone is a vertex in the tree, and possibly subtrees below this vertex
  - Each zone server knows its sub-domains’ servers
- several redundant servers per zone

# DNS records

- A: address, AAAA: IPv6 address
- NS: name server
- MX: mail exchange
- CNAME: alias
- PTR: address to name
- SRV: generic service (used for SIP)
- SOA: start of authority (gives various info. about zone)
- TXT: text

# DNS recursive queries

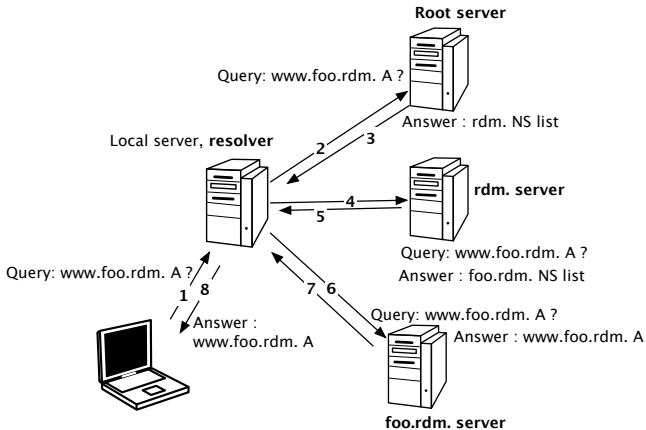
- A host asks a question to local server L (“resolver”)
- L answers if response in cache or if L is authoritative  
else
  - If L is authoritative, there is no answer
  - otherwise, L sends a request to the known server best matching the query
- D tells L the sub-servers it knows etc.

A only “talks” to L

All further traffic is between L and other servers.

L (and the request sent to L) is **recursive**, all the other servers/requests are **iterative** (they only provide a referral if not authoritative).

# DNS recursive queries (cont.)



# host & dig commands

- host (prefer to nslookup)

```
$ host www.enst-bretagne.fr
www.enst-bretagne.fr is an alias for web-srv-01-redir.enst-bretagne.fr.
web-srv-01-redir.enst-bretagne.fr has address 192.44.76.250
```

- dig

```
$dig www.enst-bretagne.fr

; <<>> DiG 9.6.0-APPLE-P2 <<>> www.enst-bretagne.fr
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 56675
;; flags: qr rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 4, ADDITIONAL: 5

;; QUESTION SECTION:
;www.enst-bretagne.fr. IN A

;; ANSWER SECTION:
www.enst-bretagne.fr. 21600 IN CNAME web-srv-01-redir.enst-bretagne.fr.
web-srv-01-redir.enst-bretagne.fr. 21600 IN A 192.44.76.250
```

## host & dig commands (cont.)

;; AUTHORITY SECTION:

```
enst-bretagne.fr. 197 IN NS rsm.enst-bretagne.fr.  
enst-bretagne.fr. 197 IN NS univers.enst-bretagne.fr.  
enst-bretagne.fr. 197 IN NS matell.enst-bretagne.fr.  
enst-bretagne.fr. 197 IN NS ns2.nic.fr.
```

;; ADDITIONAL SECTION:

```
ns2.nic.fr. 76486 IN A 192.93.0.4  
ns2.nic.fr. 76486 IN AAAA 2001:660:3005:1::1:2  
rsm.enst-bretagne.fr. 2227 IN A 192.44.77.1  
matell.enst-bretagne.fr. 197 IN A 192.108.115.2  
univers.enst-bretagne.fr. 197 IN A 192.44.75.10
```

;; Query time: 29 msec

;; SERVER: 129.88.48.1#53(129.88.48.1)

;; WHEN: Wed Mar 30 13:55:13 2011

;; MSG SIZE rcvd: 260



## host & dig commands (cont.)

```
$ dig www.imag.fr

; <<>> DiG 9.6.0-APPLE-P2 <<>> www.imag.fr
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 5463
;; flags: qr aa rd ra; QUERY: 1, ANSWER: 2, AUTHORITY: 4, ADDITIONAL: 6

;; QUESTION SECTION:
;www.imag.fr. IN A

;; ANSWER SECTION:
www.imag.fr. 172800 IN CNAME rilette.imag.fr.
rilette.imag.fr. 172800 IN A 129.88.34.211

;; AUTHORITY SECTION:
imag.fr. 172800 IN NS dns.inria.fr.
imag.fr. 172800 IN NS imag.imag.fr.
imag.fr. 172800 IN NS ns2.nic.fr.
imag.fr. 172800 IN NS isis.imag.fr.
```

## host & dig commands (cont.)

```
;; ADDITIONAL SECTION:
dns.inria.fr. 5606 IN A 193.51.208.13
ns2.nic.fr. 76312 IN A 192.93.0.4
ns2.nic.fr. 76312 IN AAAA 2001:660:3005:1::1:2
imag.imag.fr. 172800 IN A 129.88.30.1
imag.imag.fr. 172800 IN AAAA 2001:660:5301:1e::101
isis.imag.fr. 172800 IN A 129.88.32.24

;; Query time: 9 msec
;; SERVER: 129.88.48.1#53(129.88.48.1)
;; WHEN: Wed Mar 30 13:58:07 2011
;; MSG SIZE rcvd: 272
```

# Delegation system

```
$dig +trace fr.  
; <<>> DiG 9.2.1 <<>> +trace fr.  
;; global options: printcmd  
.           190181  IN      NS      C.ROOT-SERVERS.NET.  
.           190181  IN      NS      D.ROOT-SERVERS.NET.  
.           190181  IN      NS      E.ROOT-SERVERS.NET.  
.           190181  IN      NS      F.ROOT-SERVERS.NET.  
.           190181  IN      NS      G.ROOT-SERVERS.NET.  
.           190181  IN      NS      H.ROOT-SERVERS.NET.  
.           190181  IN      NS      I.ROOT-SERVERS.NET.  
.           190181  IN      NS      J.ROOT-SERVERS.NET.  
.           190181  IN      NS      K.ROOT-SERVERS.NET.  
.           190181  IN      NS      L.ROOT-SERVERS.NET.  
.           190181  IN      NS      M.ROOT-SERVERS.NET.  
.           190181  IN      NS      A.ROOT-SERVERS.NET.  
.           190181  IN      NS      B.ROOT-SERVERS.NET.  
;; Received 244 bytes from 129.88.38.1#53(129.88.38.1) in 2 ms  
  
fr.         172800  IN      NS      DNS.CS.WISC.EDU.  
fr.         172800  IN      NS      DNS.INRIA.fr.
```

## Delegation system (cont.)

```
fr.                172800  IN      NS      DNS.PRINCETON.EDU.
fr.                172800  IN      NS      NS1.NIC.fr.
fr.                172800  IN      NS      NS2.NIC.fr.
fr.                172800  IN      NS      NS3.NIC.fr.
fr.                172800  IN      NS      NS3.DOMAIN-REGISTRY.NL.
fr.                172800  IN      NS      NS-EXT.VIX.COM.
```

;; Received 351 bytes from 192.33.4.12#53(C.ROOT-SERVERS.NET) in 97 ms

```
fr.                86400  IN      SOA     ns1.nic.fr. nic.nic.fr. 2004013000 21600 3600
```

;; Received 68 bytes from 128.105.2.10#53(DNS.CS.WISC.EDU) in 131 ms

**The reply order changes from one query to the next ⇒  
load balancing**


(The number on each line is the entry TTL)

# Reverse lookup – PTR record

```
$host 129.88.38.1
```

```
1.38.88.129.in-addr.arpa domain name pointer horus.imag.fr.
```

The address bytes are reversed: : “least significant” portion first, more “general”, last



## Reverse lookup – PTR record (cont.)

```
$dig -x 129.88.55.113
```

```
; <<>> DiG 9.10.6 <<>> -x 129.88.55.113
```

```
;; global options: +cmd
```

```
;; Got answer:
```

```
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 18414
```

```
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 4
```

```
;; OPT PSEUDOSECTION:
```

```
; EDNS: version: 0, flags:; udp: 4096
```

```
;; QUESTION SECTION:
```

```
;113.55.88.129.in-addr.arpa. IN PTR
```

```
;; ANSWER SECTION:
```

```
113.55.88.129.in-addr.arpa. 3600 IN PTR molene.imag.fr.
```

```
;; AUTHORITY SECTION:
```

```
55.88.129.in-addr.arpa. 3600 IN NS eip1.u-ga.fr.
```

```
55.88.129.in-addr.arpa. 3600 IN NS eip2.u-ga.fr.
```

## Reverse lookup – PTR record (cont.)

;; ADDITIONAL SECTION:

eip1.u-ga.fr. 488 IN A 195.83.24.250

eip2.u-ga.fr. 488 IN A 152.77.1.250

eip2.u-ga.fr. 2870 IN AAAA 2001:660:5301:1e::101

;; Query time: 58 msec

;; SERVER: 130.190.190.2#53(130.190.190.2)

;; WHEN: Wed Mar 20 21:33:30 CET 2019

;; MSG SIZE rcvd: 186

# Name Resolution in C

- IPv4 name query in C : `gethostbyname()`
- Better method (v4 & v6 compliant) :

```
hints.ai_flags          = 0;
hints.ai_family         = PF_UNSPEC; // the result can be IPv4 or IPv6
hints.ai_socktype       = SOCK_STREAM;
hints.ai_protocol       = IPPROTO_TCP;
hints.ai_addrlen        = 0;
hints.ai_addr           = NULL;
hints.ai_canonname      = NULL;
hints.ai_next           = NULL;
...
getaddrinfo(hostname, portnum, &hints, &rai)
//Create the socket, using the protocol of the first answer
s = socket(rai->ai_family, rai->ai_socktype, rai->ai_protocol);
// connect to the address in the first answer
connect(s, rai->ai_addr, rai->ai_addrlen);
```

(This code does not include the usual error traps, and could loop through the answers cf. `man getaddrinfo`)