Plain "Old" DNS

WACREN, DNS/DNSSEC Regional Workshop

Ouagadougou, 10-14 October 2016

IP: Identifiers on the Internet

- The fundamental identifier on the internet is an IP address.
- Each host connected to the Internet has a unique IP address
  - IPv4 or IPv6
  - Uniqueness guaranteed through allocation from one single pool

How Devices use Identifiers

- On operating system level only the numbers matter
- Terminology in this context
  - TCP/IP Stack
  - Sockets
- The devices do not care about names

What is easier to remember?

- Humans tend to remember names better, easier to associate

NL 1098VA 419	or	Kruislaan 419, Amsterdam, Netherlands
TG 9613 AL		or Alain Hyundai X35
178.79.184.95	or	www.wacren.net

host.txt

- In the1970's ARPA net, tables where maintained mapping host-names to IP addresses
  - SRI-NIC
  - Tables were pulled from the single machine
  - Problems
    - traffic and load
    - Name collisions
    - Consistency

DNS

- Domain Name System provides a scalable, distributed lookup mechanism.
- DNS created in 1983 by Paul Mockapetris
  - RFCs 882 and 883
- IETF Full Standard: RFCs 1034 and 1035 (1987)
  - modified, updated, and enhanced
  - DNS Security extensions being the most recent

The four components

- A “name space”
- Servers making that name space available
- Resolvers (clients) which query the servers about the name space
- The protocol
  - Glues all together

The Namespace Design

- The namespace needs to be made hierarchical to be able to scale
  - Both “technical” and “managerial” delegation
  - Control of parts of the namespace follows the hierarchy
  - Hierarchy represented in labels

    `country.nren.wacren.net`

The namespace: Domains

- Domains are “namespace subsets”
- Everything below .com is in the com domain.
- Everything below foo.net is in the foo.net domain and in the net domain.

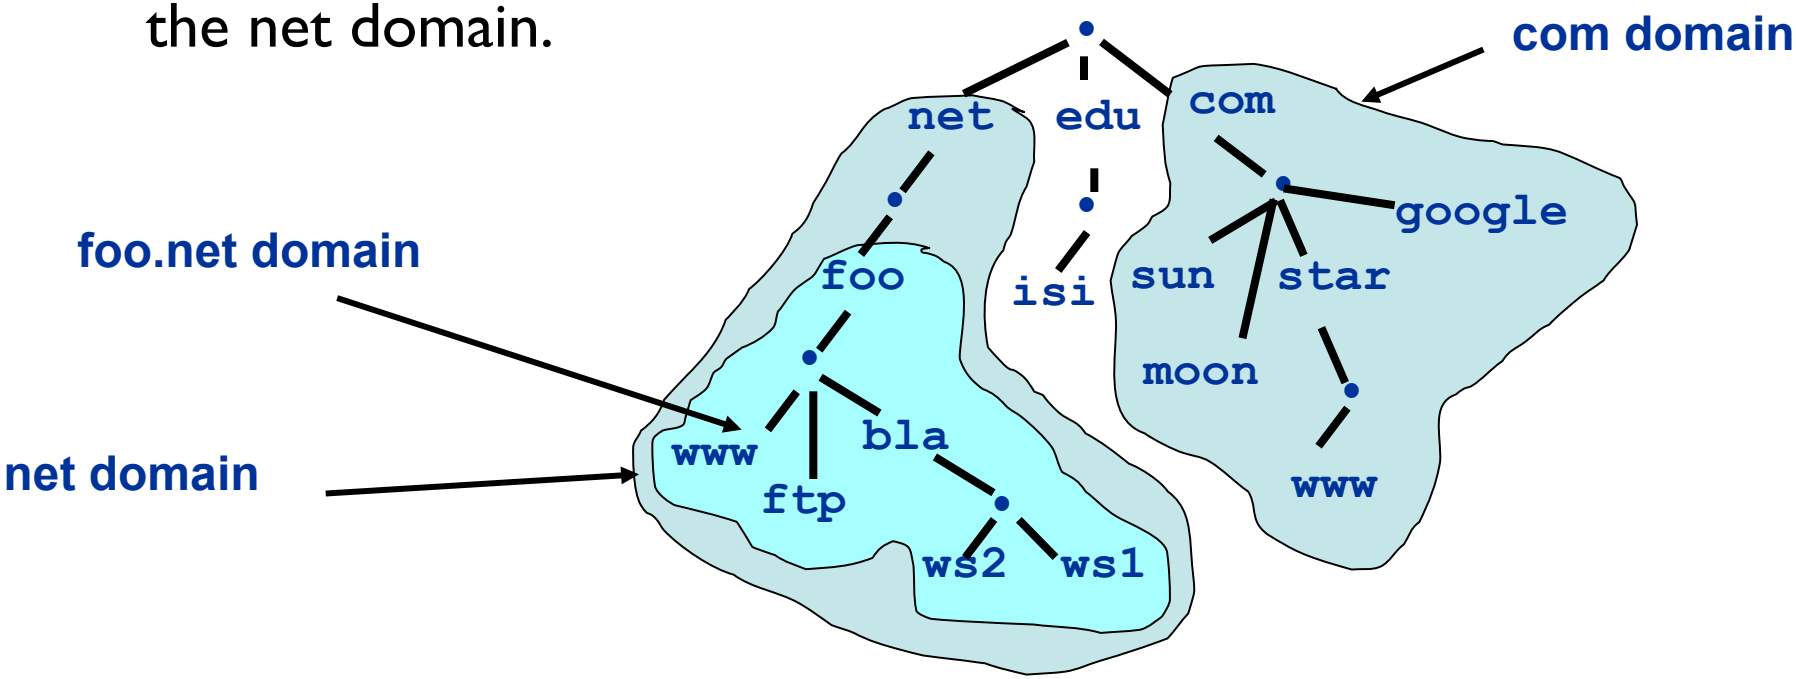

The namespace: Zones and Delegations

- Zones are “administrative spaces”
- Zone administrators are responsible for portion of a domain’s name space
- Authority is delegated from a parent and to a child

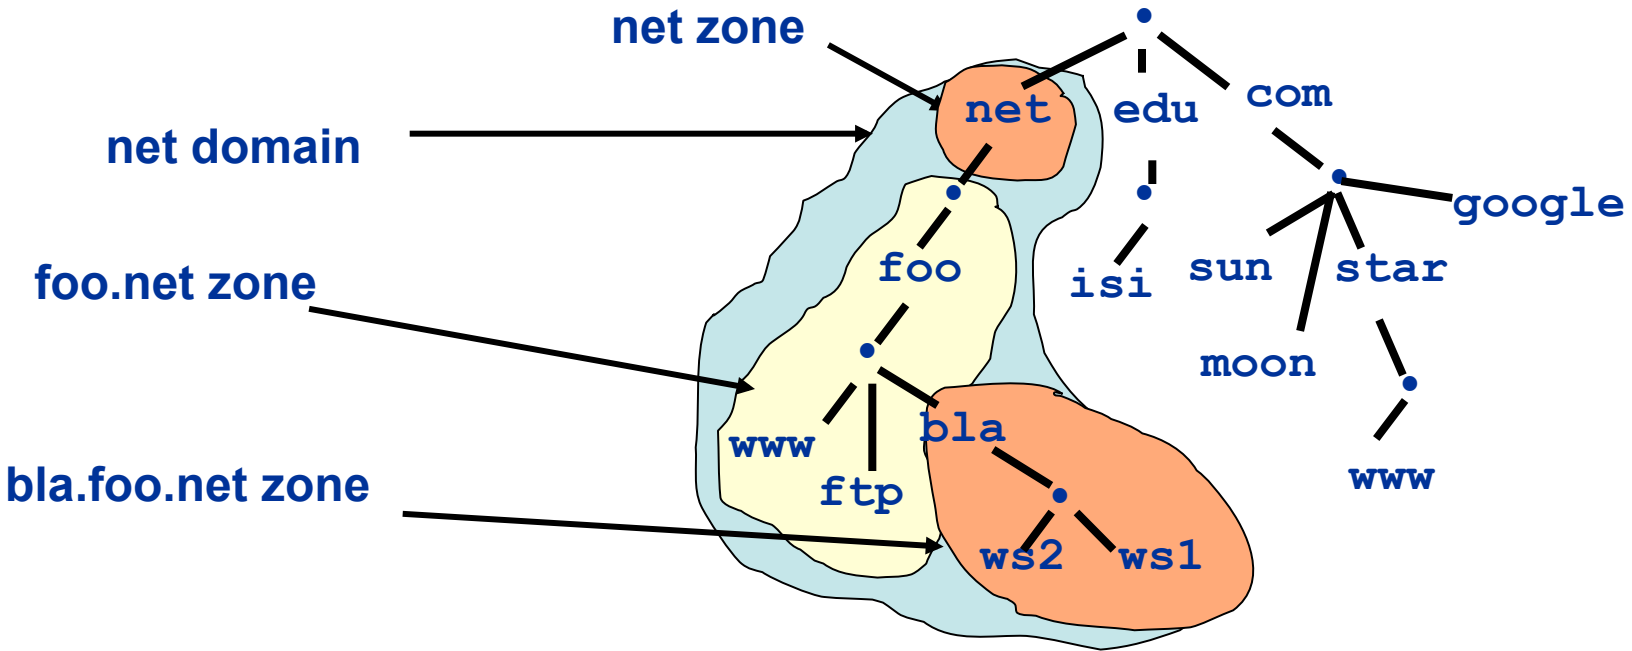

Some Jargon

Top-Level Domains (TLD)

Country cctld
Generic gtld

Second-Level Domains

In practice TLDs
And SLDs are actually zones

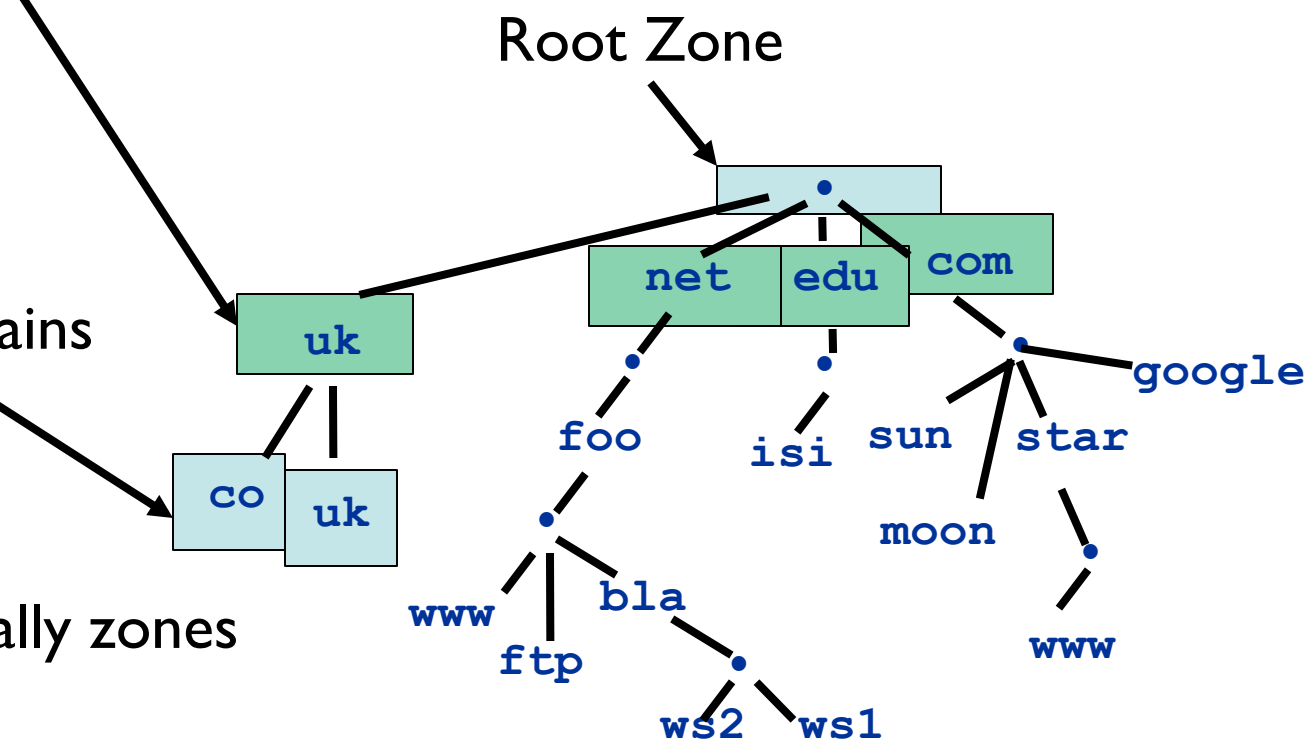

Name Servers

- Name servers answer 'DNS' questions.
- Several types of name servers
  - Authoritative servers
    - Serves the authoritative data for 'Zones'
    - Primary and Secondary
  - (Caching) recursive servers
    - Also called caching forwarders
  - Mixture of functionality

Zones are served by authoritative name servers

Each zone served by multiple servers (over 10^6) in total

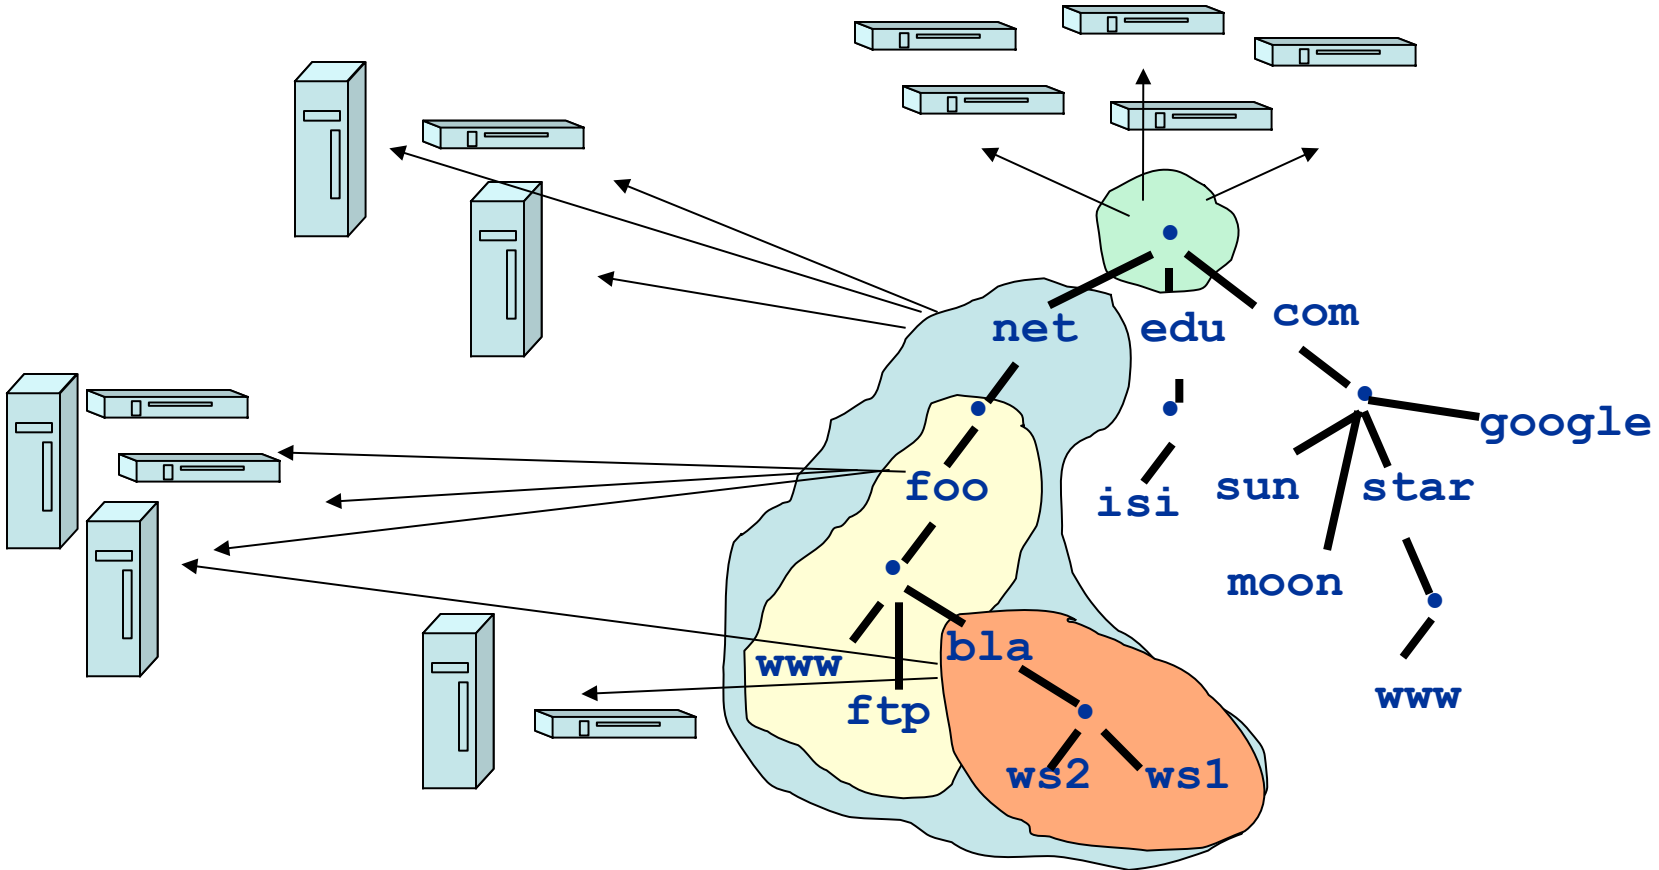

Concept: Resolving process & Cache

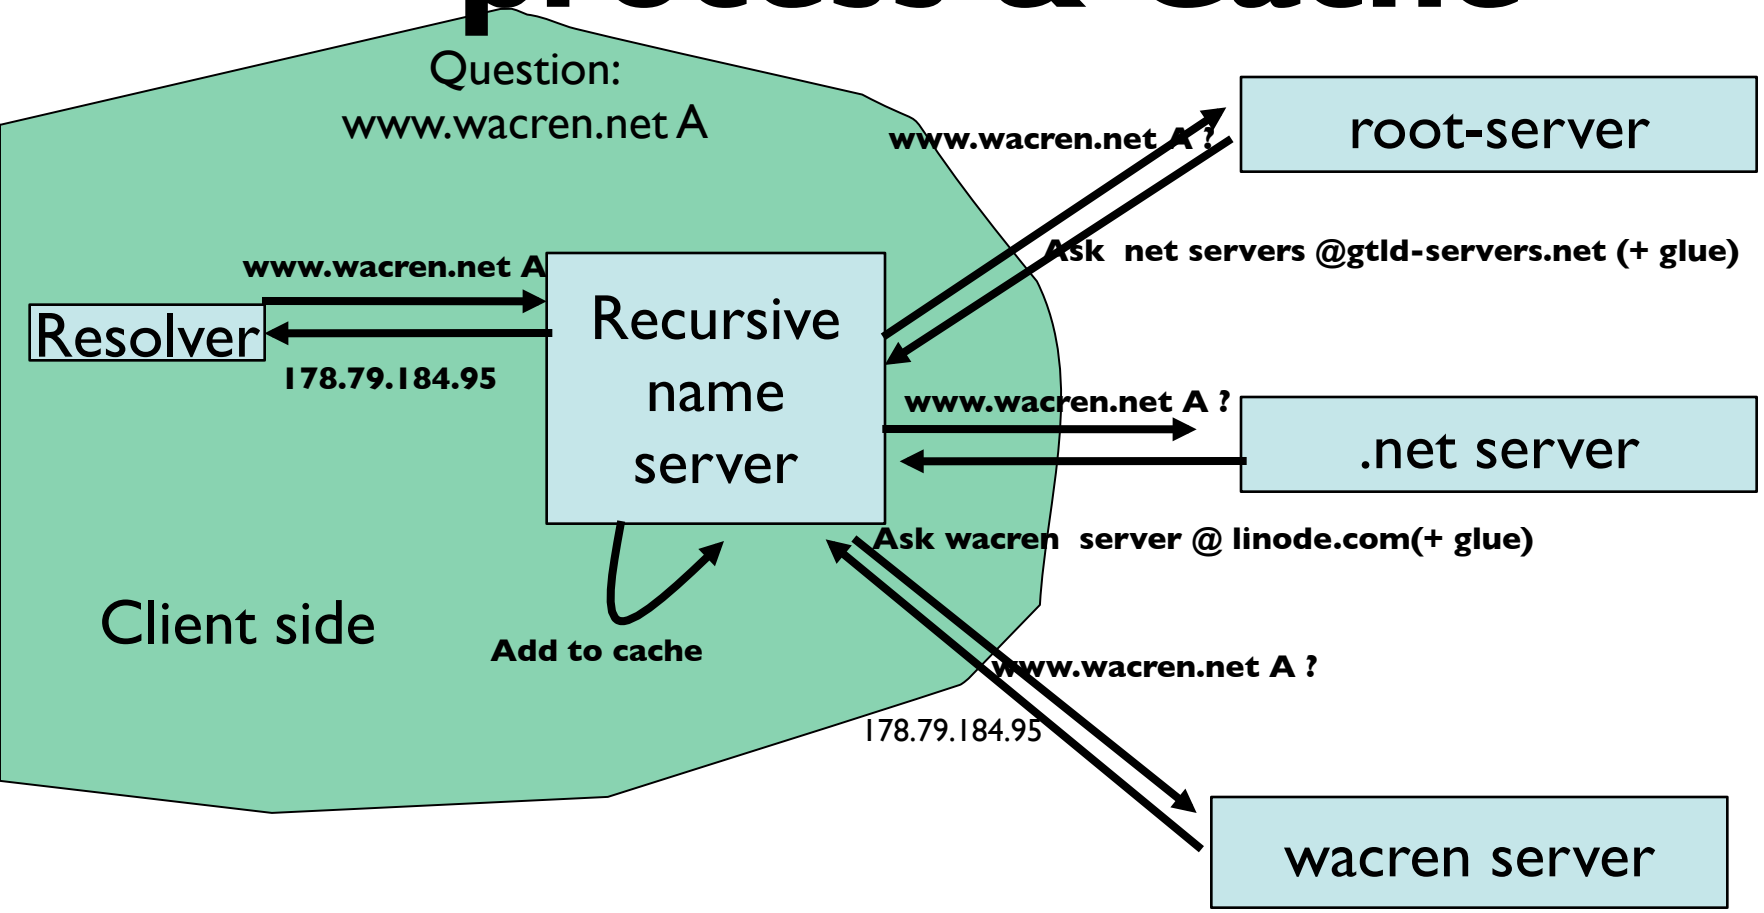

Hooking this together

Changes in DNS do not propagate instantly!

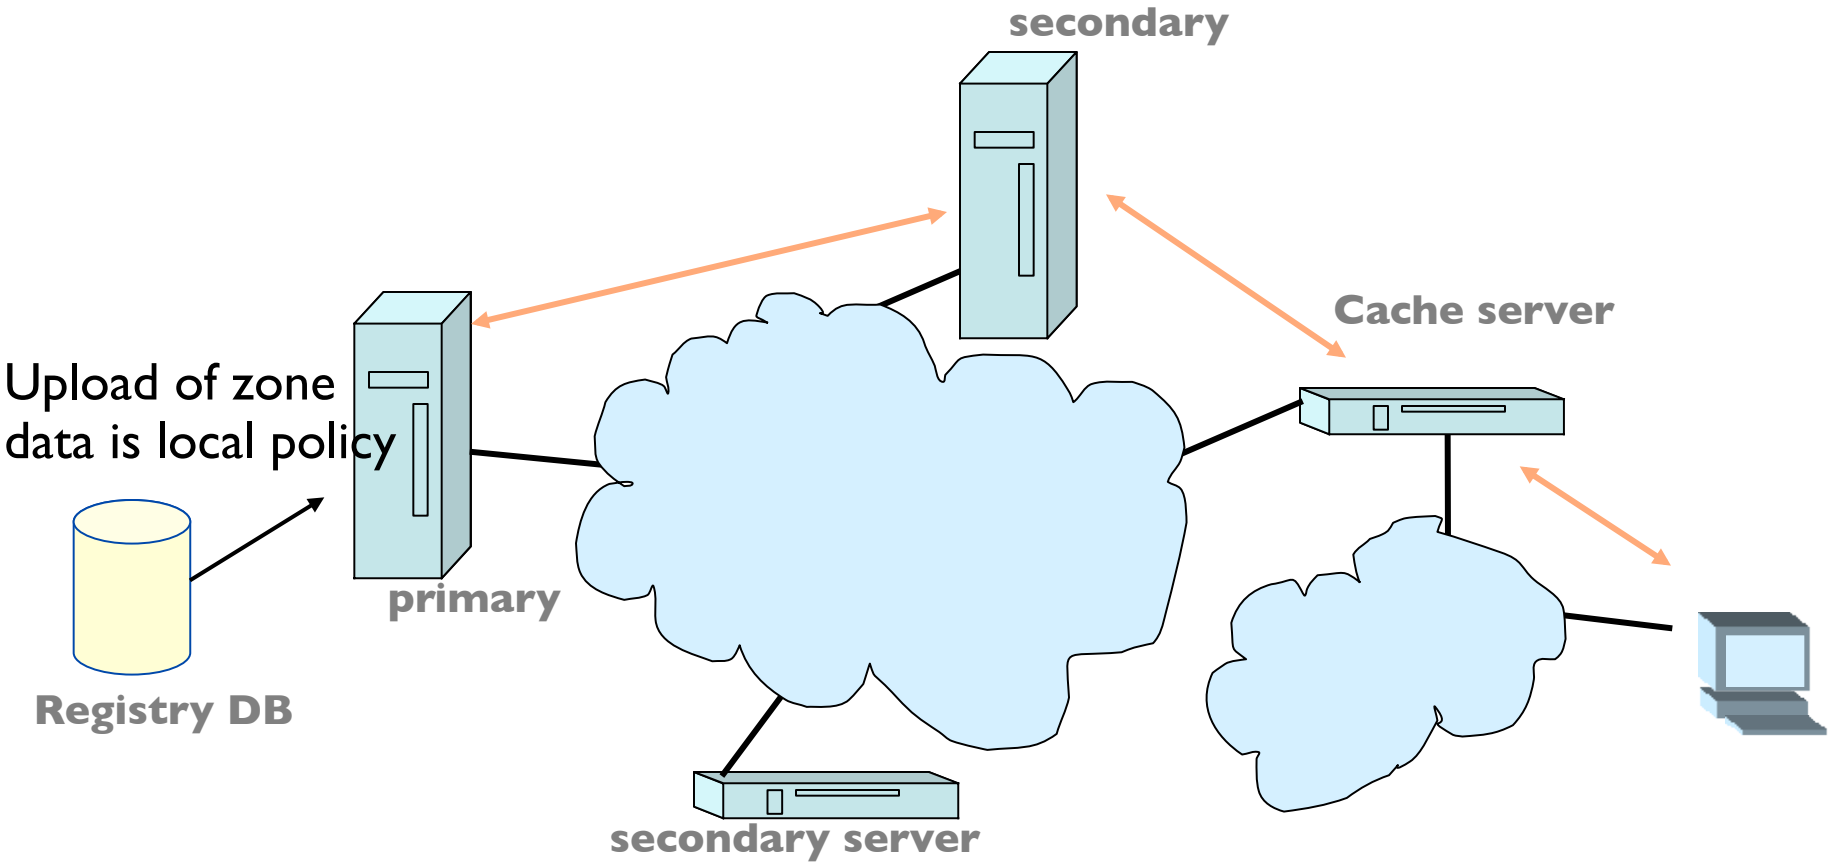

DNS Features

- A lookup mechanism for translating objects into other objects
- A globally distributed, loosely coherent, scalable, reliable, dynamic database
- Comprised of four components
  - A "name space"
  - Servers making that name space available
  - Resolvers (clients) which query the servers about the name space
  - The DNS protocol

DNS Features: Global Distribution

- Data is maintained locally, but retrievable globally
  - No single computer has all DNS data
  - Total number of servers: in the 10^6 to 10^7 range
- DNS lookups can be performed by any device
- Remote DNS data is locally cachable to improve performance

DNS Features: Loose Coherency

- The database is always internally consistent
  - Each version of a subset of the database (a zone) has a serial number
    - The serial number is incremented on each database change
- Changes to the master copy of the database are replicated according to timing set by the zone administrator
- Cached data expires according to timeout set by zone administrator
- Response the same regardless of who the source of the query

DNS Features: Scalability

- No limit to the size of the database
  - One server has over 40,000,000 names
- No limit to the number of queries
  - 24,000 queries per second handled easily by one server
- Queries distributed among primary, secondary, and caches servers

DNS Features: Reliability

- Data is replicated
  - Data from primary is copied to multiple secondaries
  - The system can deal with outage of servers
- Clients can query
  - All authoritative servers
  - No difference between primaries and secondaries
- Clients will typically query local caches
- DNS protocols can use either UDP or TCP
  - If UDP, DNS protocol handles retransmission, sequencing, etc.

DNS Features: Dynamicity

- Database can be updated dynamically
  - Add/delete/modify of any record
  - Within seconds possible, traditionally lower update rates
- Modification of the primary database triggers replication
  - Only primary can be dynamically updated

RRs and RRSets

- Resource Record:
  - name TTL class type rdata

```
www.example.com      7200   IN     A      192.0.2.3
```

- RRset: RRs with same name, class and type:

```
www.example.com      7200 IN     A      192.0.2.3
                                 A      198.51.100.3
                                 A      203.0.113.3
```

DNS Packet

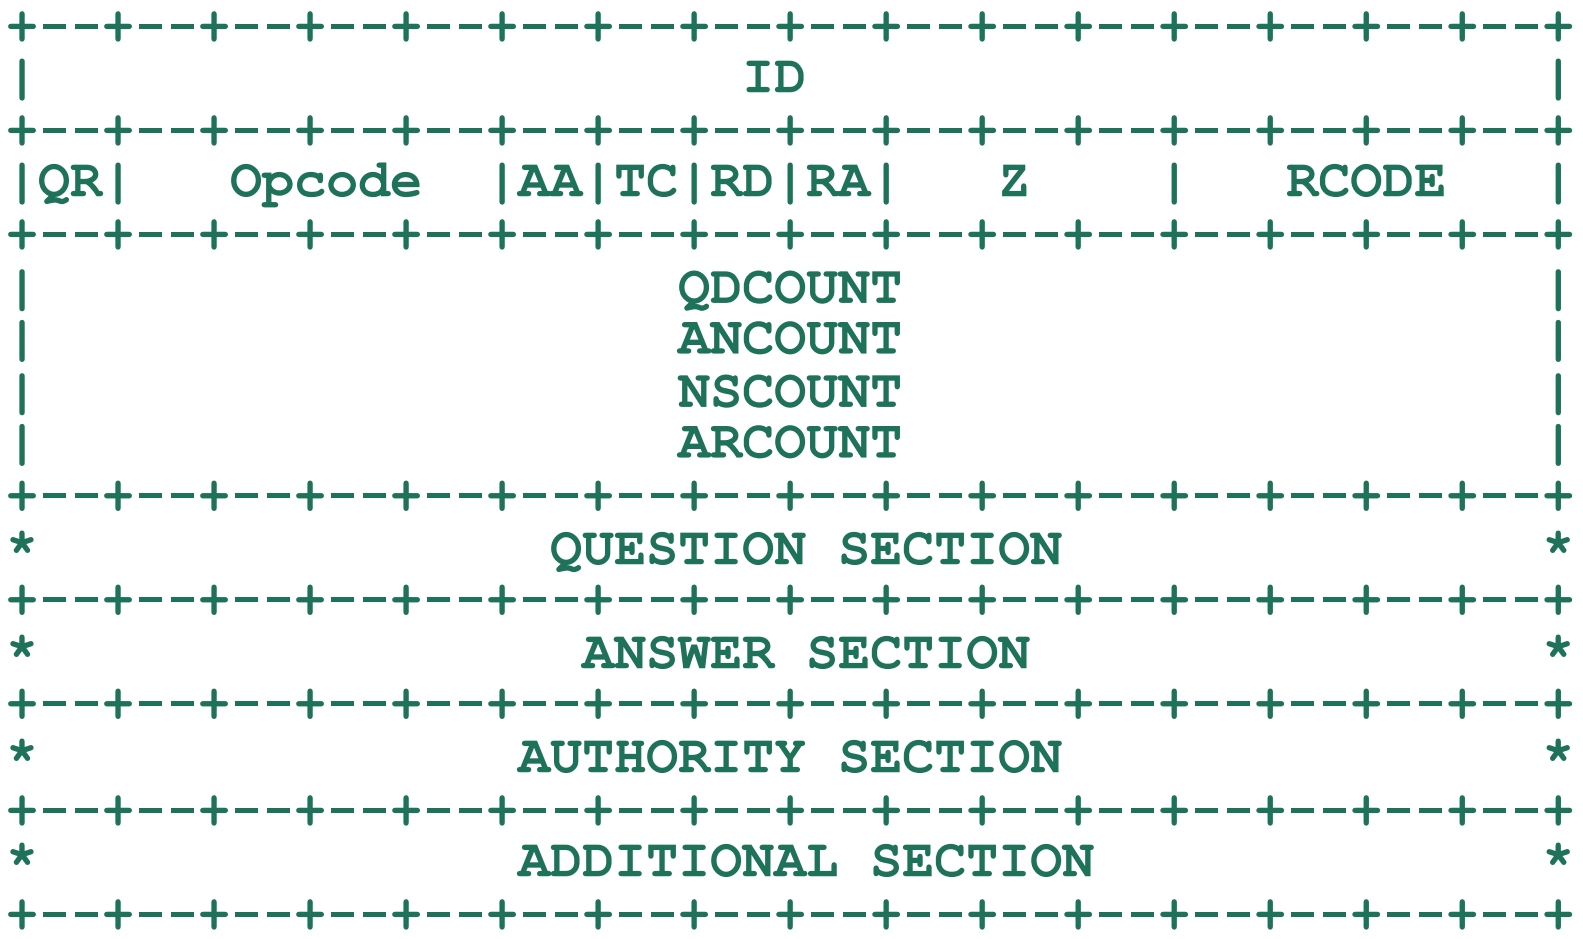
ID
QR
Opcode
AA
TC
RD
RA
Z
RCODE
QDCOUNT
ANCOUNT
NSCOUNT
ARCOUNT
QUESTION SECTION
ANSWER SECTION
AUTHORITY SECTION
ADDITIONAL SECTION

DIG and the Packet

```
; <<>> DiG 9.10.0-P2 <<>> www.wacren.net
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 2652
;; flags: qr rd ra; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 1

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 4096
;; QUESTION SECTION:
;www.wacren.net.                        IN      A

;; ANSWER SECTION:
www.wacren.net.                 83748   IN      A       178.79.184.95

;; AUTHORITY SECTION:
wacren.net.             170146  IN      NS      ns1.linode.com.
wacren.net.             170146  IN      NS      ns2.linode.com.

;; Query time: 7 msec
;; SERVER: 10.10.0.2#53(10.10.0.2)
;; WHEN: Wed Sep 28 20:38:53 MUT 2016
;; MSG SIZE  rcvd: 105
```

REVERSE DNS

WHY

- Whom clients/users are ?

- Every DNS entry name-IP (A record) must have a correspondence IP-name(PTR record)
- Otherwise:
  - Acces denied to certains services (ftp, mail, IRC,....)
  - Hard network debug (traceroute)
  - More undesirable network traffic

-

HOW

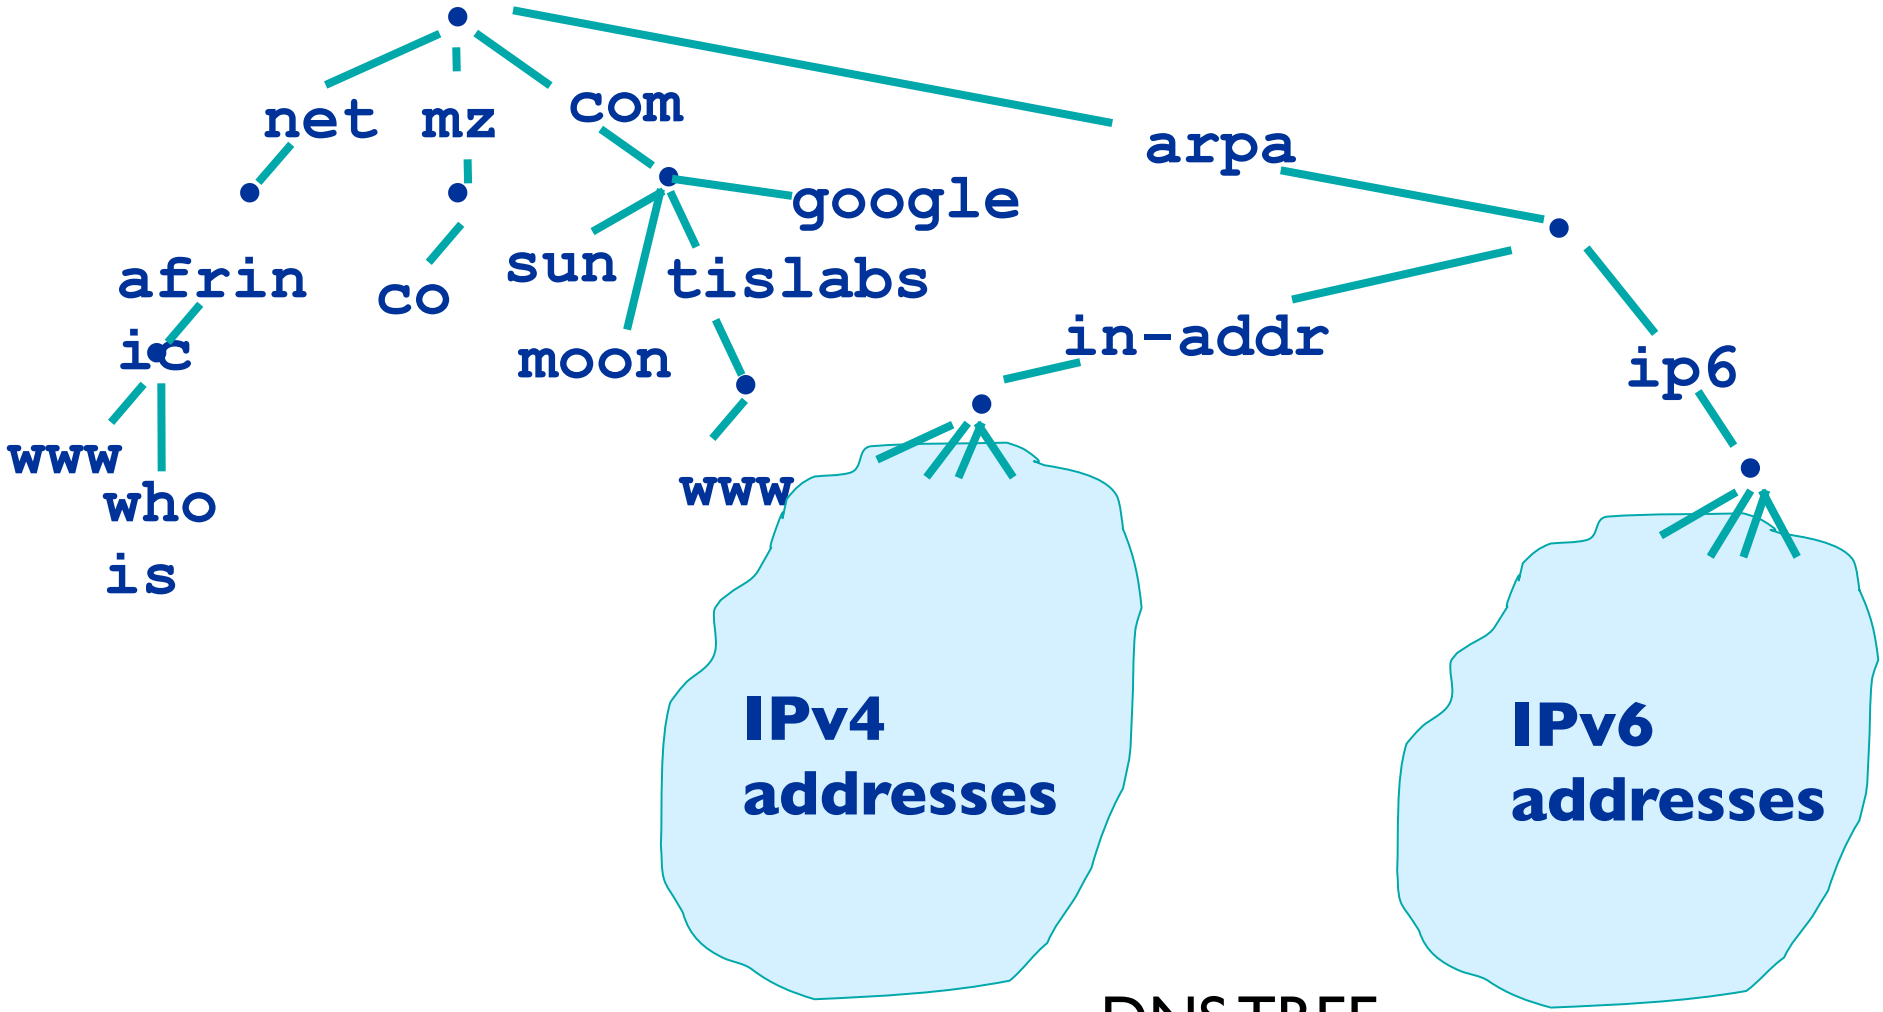
net
mz
com
arpa
google
afrin
ic
co
sun
tislabs
moon
in-addr
ip6
www
who
is
www
IPv4
addresses
IPv6
addresses

DNS TREE

IPV4

Mapping IPv4 address in DNS

- Example 196.26.1.3
  - 196/8 is allocated to RIR
  - 196.26/16 is allocated by RIR to LIR/ISP
  - 192.26.1/24 is assigned by ISP to a company.
- Delegation in the DNS:
  - in-addr.arpa delegates 196 domain to RIR
  - RIR delegates "26" sub-zone to ISP
  - ISP delegates "1" sub-zone to company.
- Name that makes this possible:
  - 1.26.192.in-addr.arpa.

Mapping IPv4 address to names

- In IPv4 the mapping is done on 8 bit boundaries(class full), address allocation is class less
  - /8, /16, /24
- Zone administration does not always overlap address administration
- If you have a /22 of address space: divide it in /24s and request a delegation for each one of them

LIR and end-users PI

- Configure your authoritative NS for the reverse zones
  - Follow DNS recommendations (RFC 2182,1912)
- Create the **`domain`** object in the RIR database
  - Only /16 and /24
- If authentication and dns check are OK, delegation is visible next time RIR push zone file

End-users

- Configure your authoritative NS for the reverse zones
  - Follow DNS recommendations (RFC 2182,1912)
- Contact your ISP
  - >=/24
- For < /24
  - RFC 2317

domain object

- domain: 209.32.196.in-addr.arpa
- descr: ubuntunet allocation
- nserver: disa.tenet.ac.za
- nserver: v6rev.tenet.ac.za
- org: ORG-UAFR1-AFRINIC
- admin-c: RJ1-AFRINIC
- tech-c: AA28-AFRINIC
- tech-c: RJ1-AFRINIC
- zone-c: AA28-AFRINIC
- mnt-by: ubunt-mnt
- mnt-lower: ubunt-mnt
- remarks: www.ubuntunet.net
- source: AFRINIC # Filtered

IPV6

Allocations policy

- Allocations policy
  - /12 allocated to RIR
  - /32 allocated to LIR/ISP
  - /48 assigned to end users in general
  - /64 assigned to end users when only one net is used
  - /128 assigned to end users when only one device is used
- Policy is moving

Mapping IPv6 address in DNS

- Number is translated into 4 bit nibbles under the ip6.arpa.

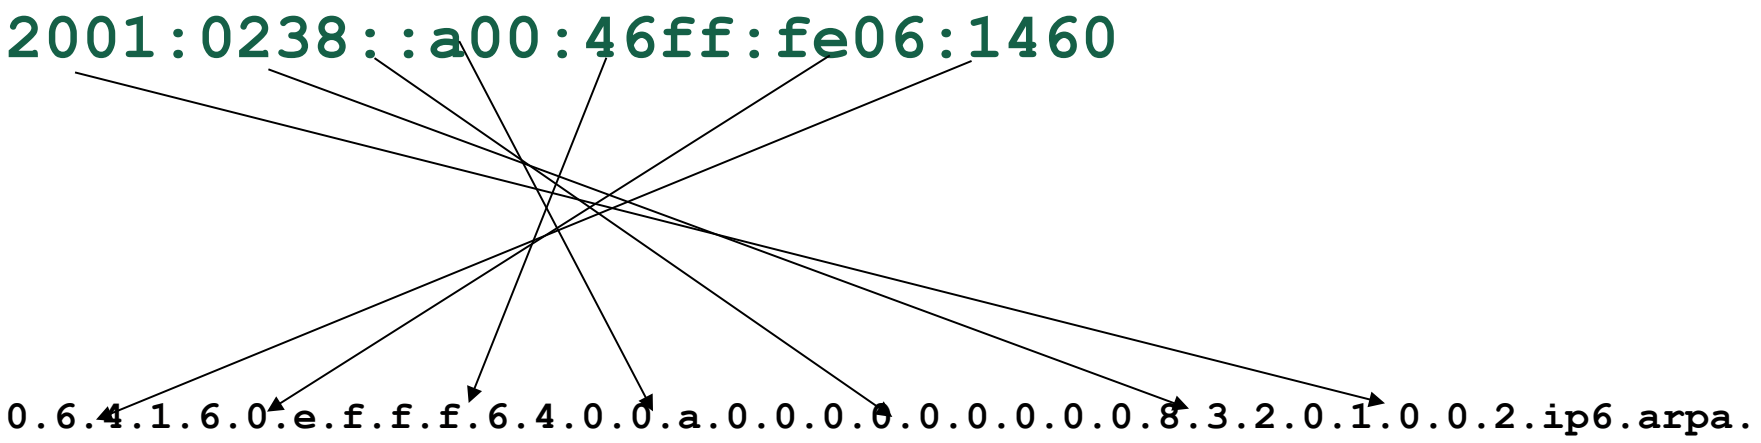

If you have a /32, split into 2 /32s

If you have a /47, split into 2 /48s

LIR and end-users PI

- Configure your authoritative NS for the reverse zones
  - Follow DNS recommendations (RFC 2182,1912)
- Create the `domain` object in the RIR database
  - /32, /48
- If authentication and dns check are OK, delegation is visible next time RIR pushes zone file

End-users

- Configure your authoritative NS for the reverse zones
  - Follow DNS recommendations (RFC 2182,1912)
- Contact your ISP
  - /48, /64, /128

domain object

```
domain:         0.6.6.0.1.0.0.2.ip6.arpa
descr:          Reverse delegation for Renater sub-TLA
admin-c:        BT261-RIPE
tech-c:         BT261-RIPE
tech-c:         GR1378-RIPE
zone-c:         GR1378-RIPE
nserver:        ns1.renater.fr
nserver:        imag.imag.fr
nserver:        ns3.nic.fr
nserver:        ns2.renater.fr
mnt-by:         RENATER-MNT
remarks:        changed:        rensvp@renater.fr 20021112
remarks:        changed:        rensvp@renater.fr 20100527
created:        2002-11-12T14:14:47Z
last-modified:  2015-08-07T13:30:20Z
source:         RIPE
```

QUESTIONS?