# DNS : Domain Name System

*Surasak Sanguanpong*
*nguan@ku.ac.th*
*http://www.cpe.ku.ac.th/~nguan*

Last updated: May 24, 1999

# Outline

- DNS basic
- name space
- name resolution process
- protocol
- configurations

# Why need DNS?

- host table /etc/hosts compiled from *HOST.TXT* (maintain by SRI NIC)
  - simple text file with has IP address to name mapping
- problems
  - traffic and load
  - name collision
  - consistency
- A hierarchical name with distributed control is needed

# DNS basic

- DNS is a distributed database
- TCP/IP applications use DNS to
  - map hostname to IP address
  - map IP address to hostname
  - provide e-mail routing information
    - mail nguan@cpe.ku.ac.th => mail.cpe.ku.ac.th
  - handle aliases
    - www.eng.ku.ac.th is actually is1.eng.ku.ac.th

# Naming Scheme

**cc.cpe.ku.ac.th**

**more specific**

**building tree from top to bottom**

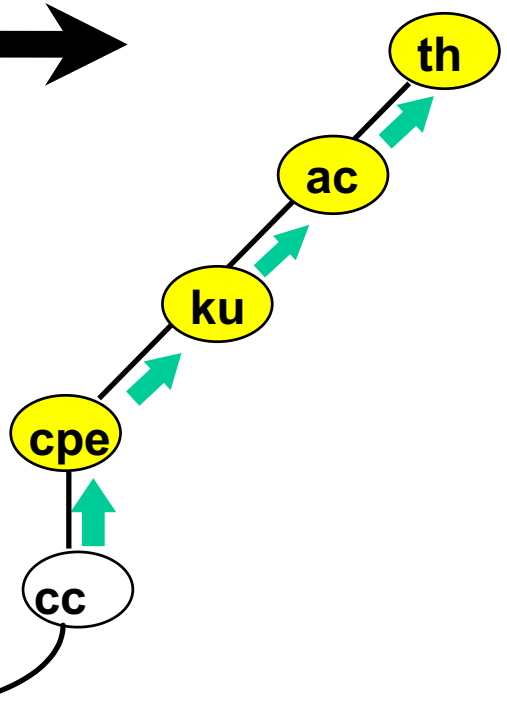


- **name space is a tree of domain**
- **names are case-insensitive**

# Domain Name Space

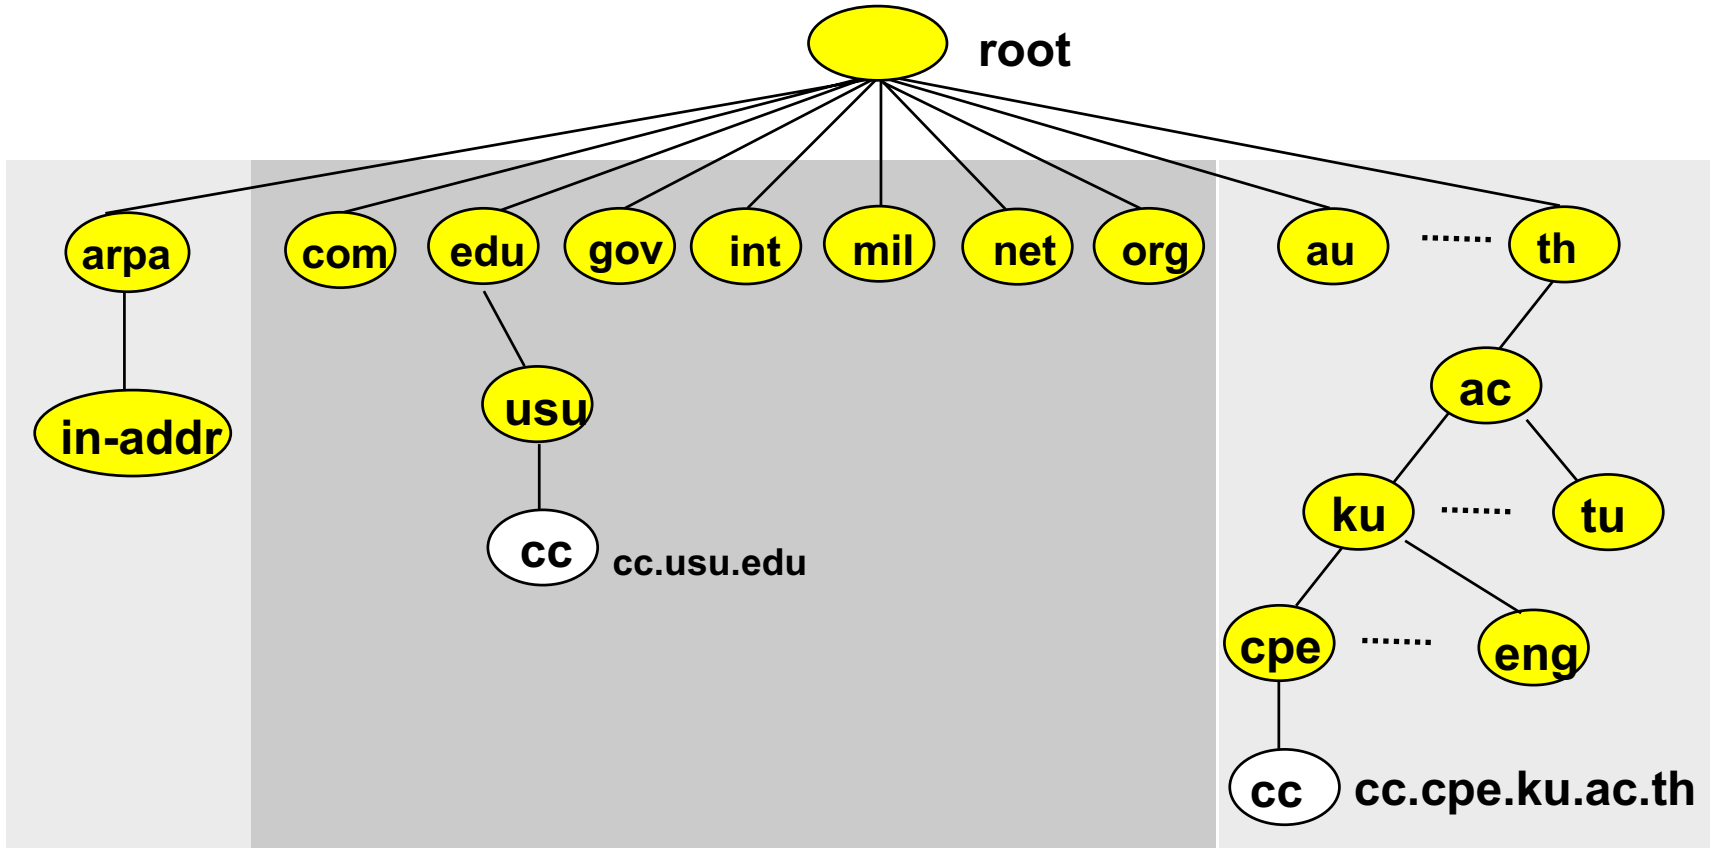

# DNS Management

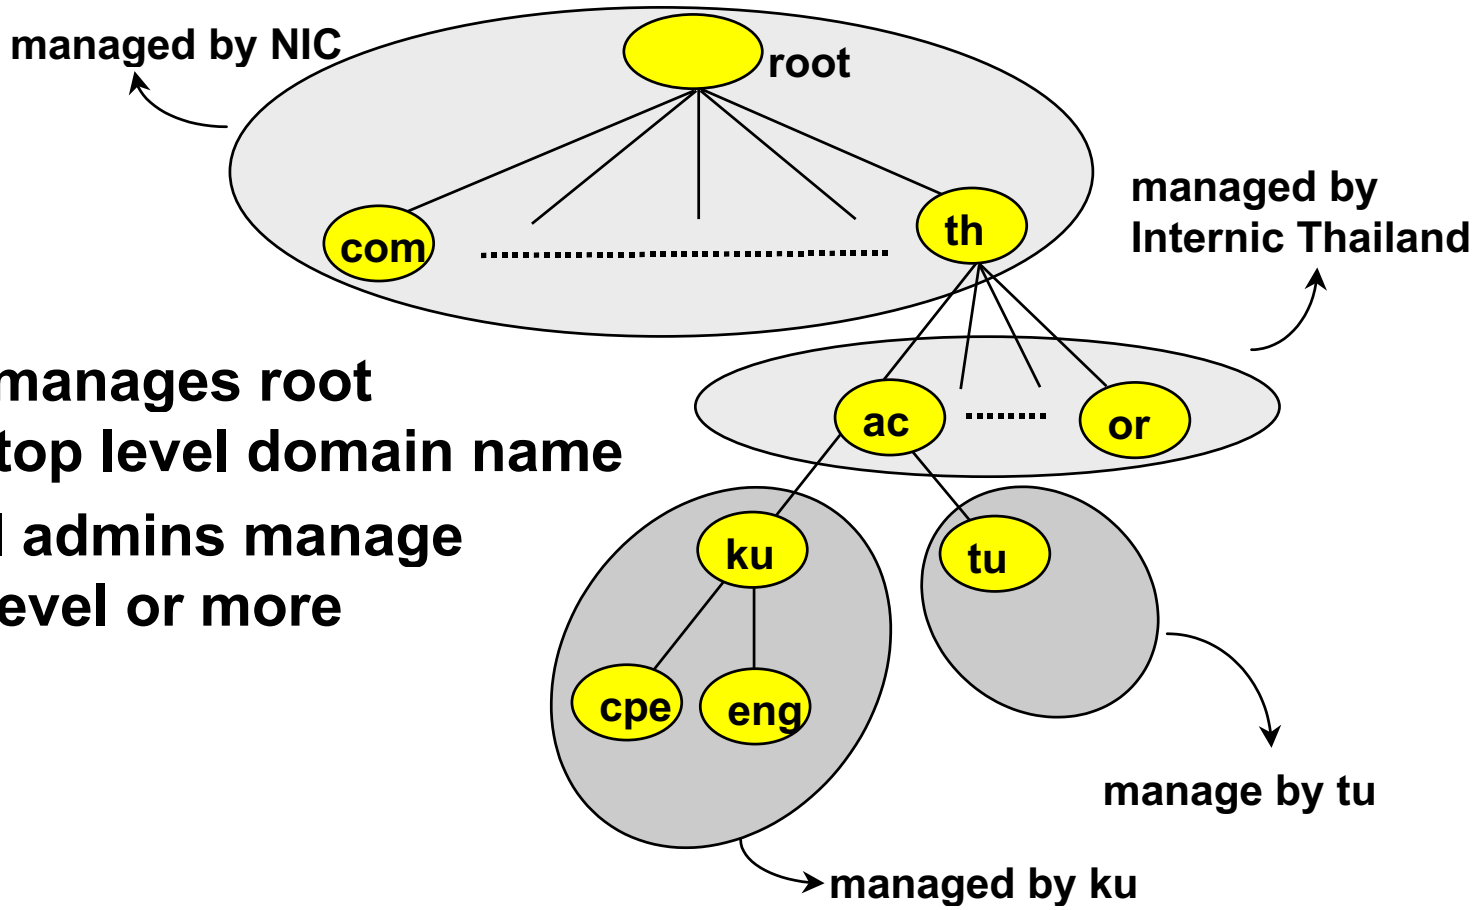


- NIC manages root and top level domain name
- local admins manage 3rd level or more

# Domain Name Concept

***label*** every node has a label (except root)

***domain name*** the list of labels, starting at that node, working up to the root, using a "." to separate
*e.g. ku.ac.th, cpe.ku.ac.th*

***absolute domain name*** domain name that ends with a period
*e.g. cc.cpe.ku.ac.th.*

***relative domain name*** name to be completed
*e.g. cc*

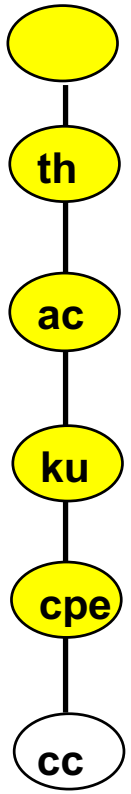

# Domains

*domain* subtree of the domain name space

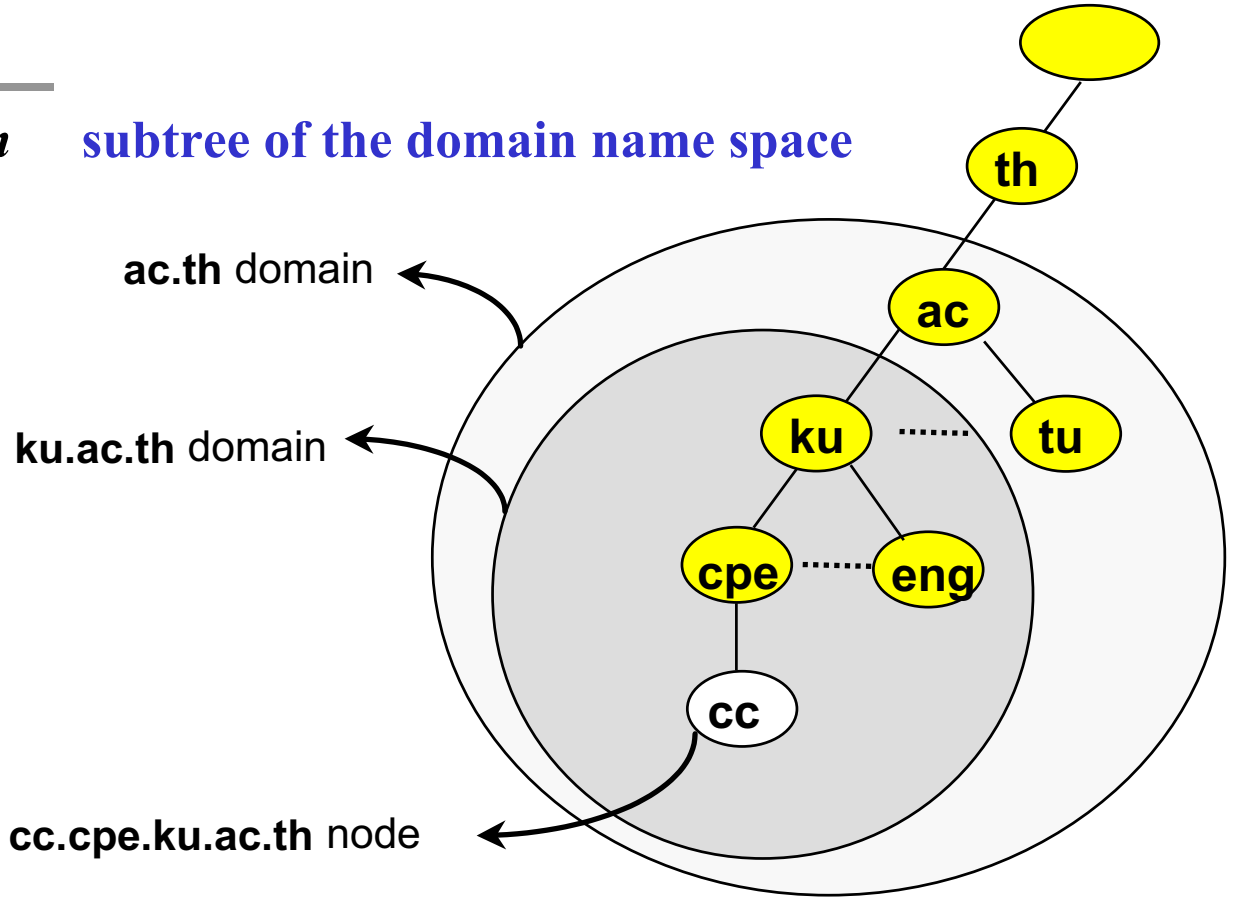

# Domains and Zones

- ***Zone*** **is a subtree for which naming authority has been *delegated***

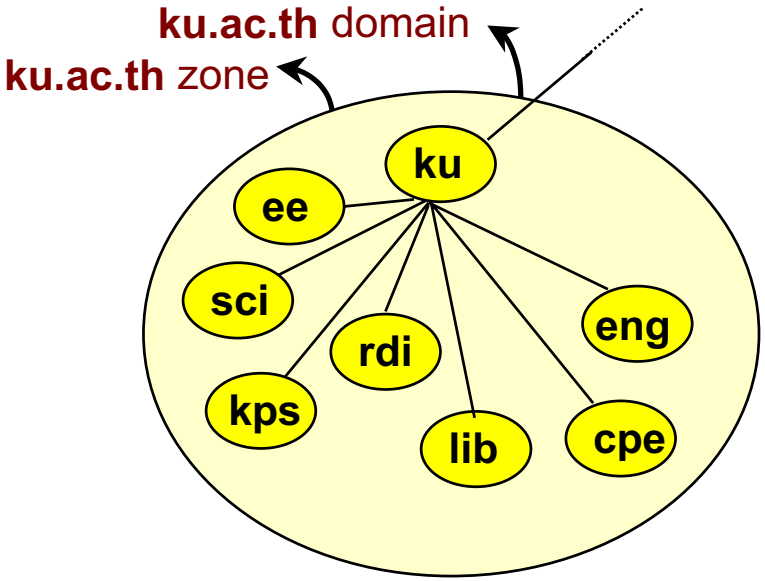


**case 1 : single DNS administration**

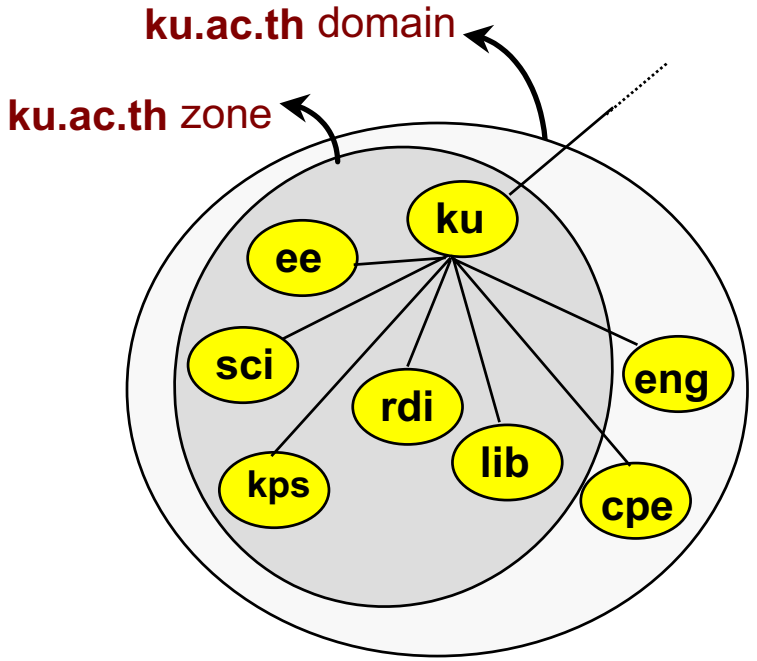


**case 2 : cpe and eng have authority for their zones**

# Name Servers

- ***Name server*** **: Server that store information about the zone**

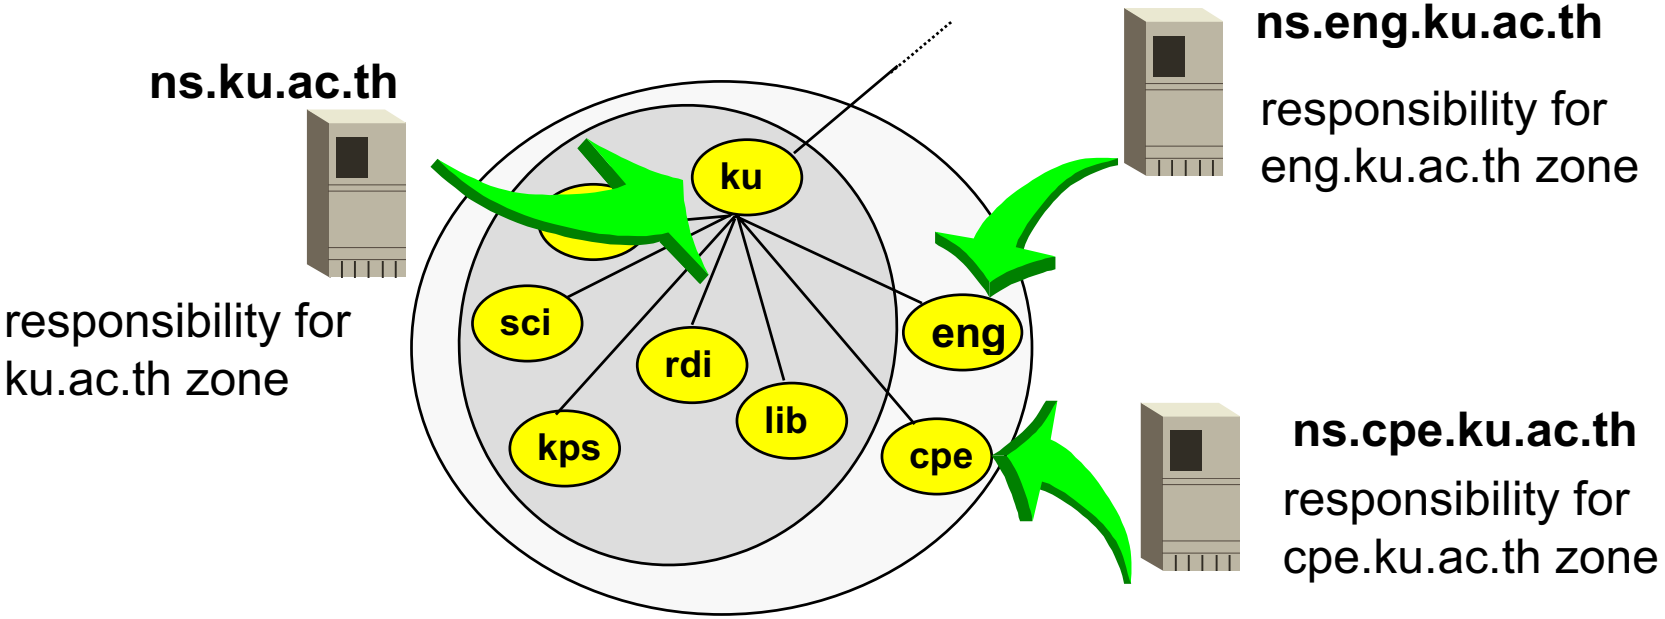

# Type of Name Servers

- *Primary Name server* gets the data for zones from files on the host it runs on
- *Secondary Name server* gets its zone data from the primary for *redundancy* and *workload distribution*

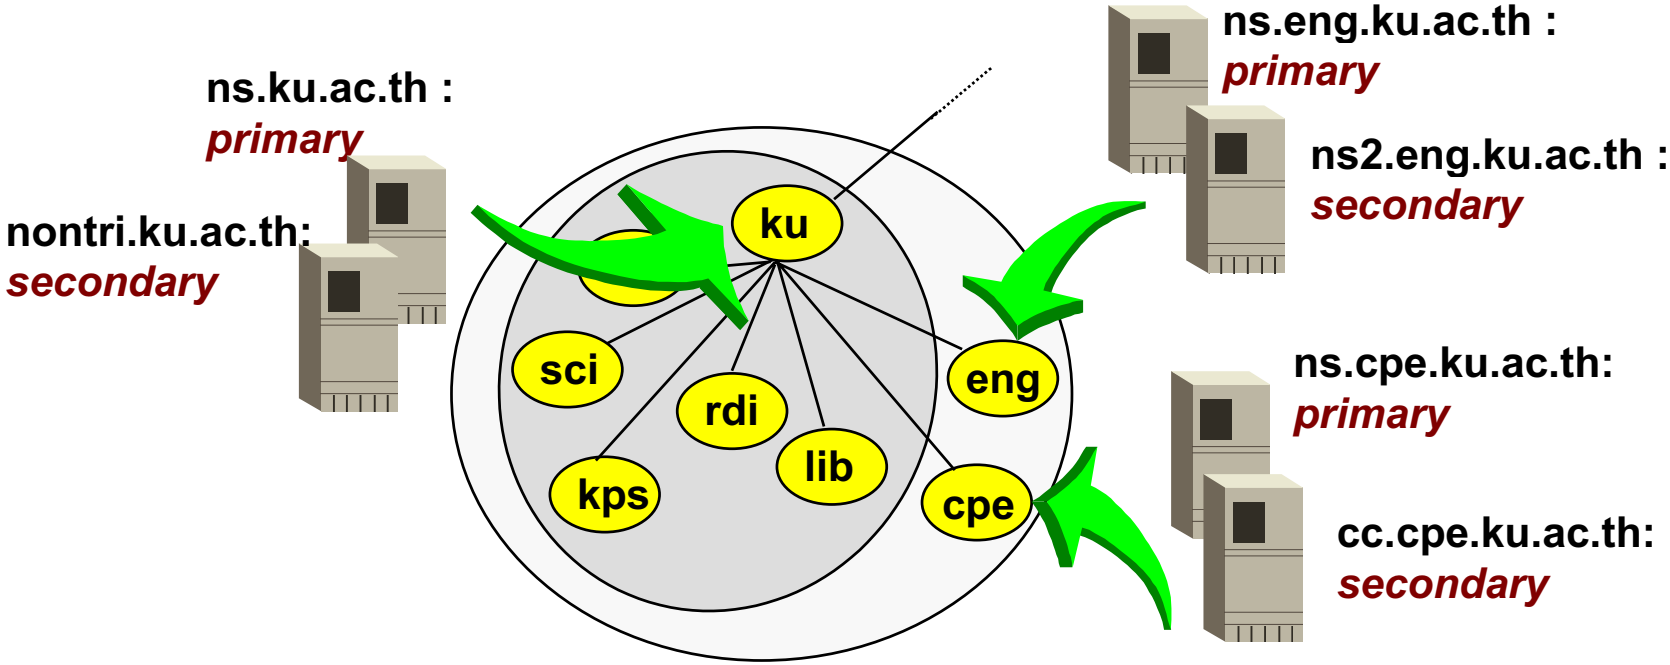

# Zone Transfer

- ***Secondary Name Server pulls zone data over from the primary* called *zone transfer.***

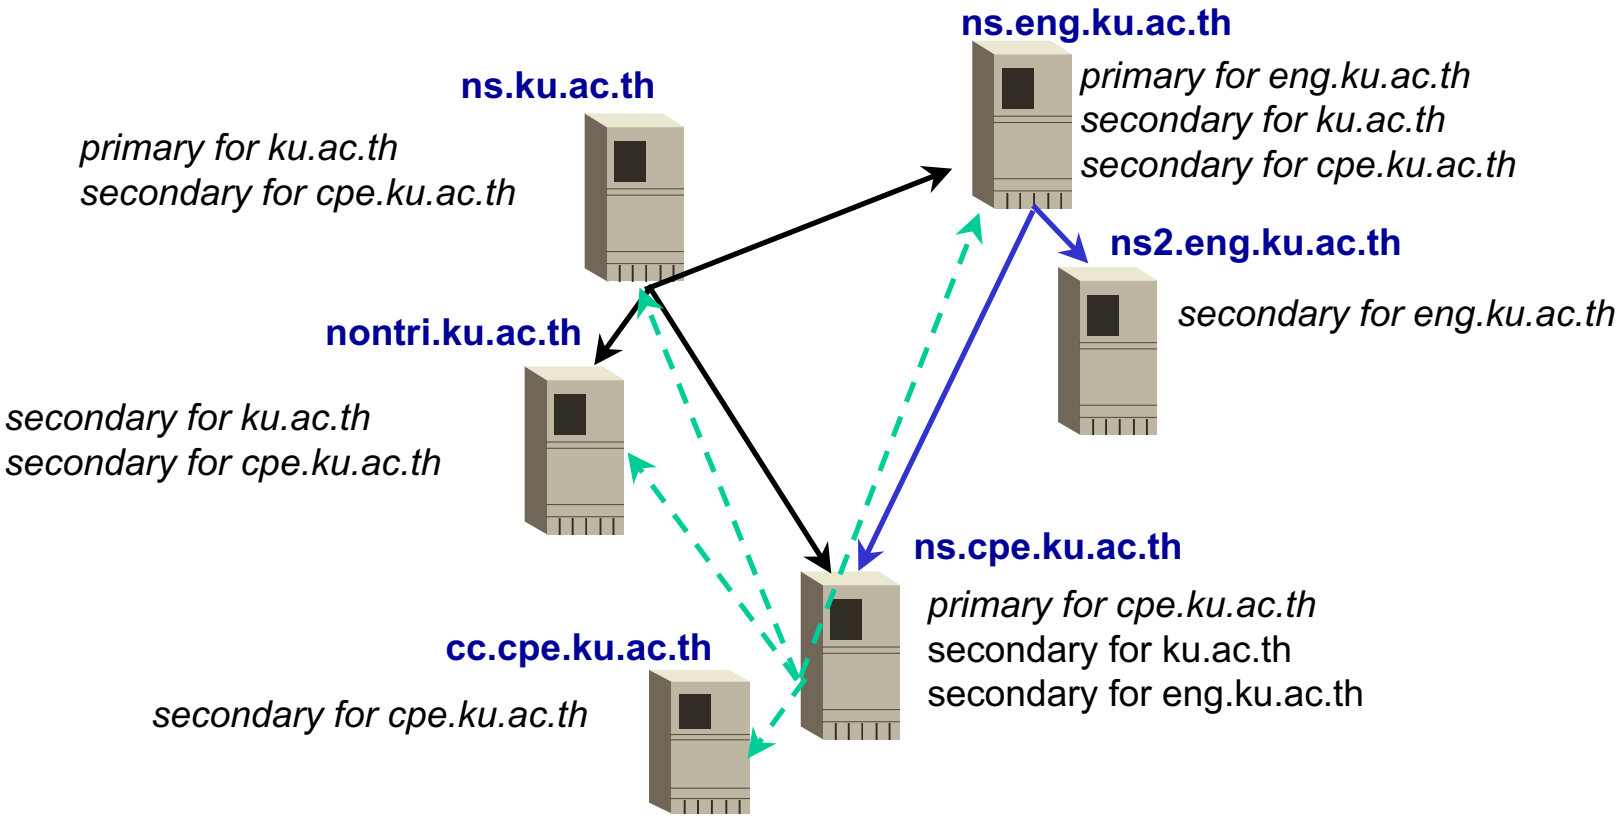

# Root Name Server

- name server must contact other name servers for non local IP
- it has to know IP address of the top most server called *root name server*
- root name server - provide the names and address of the name server authoritative for top level domain name

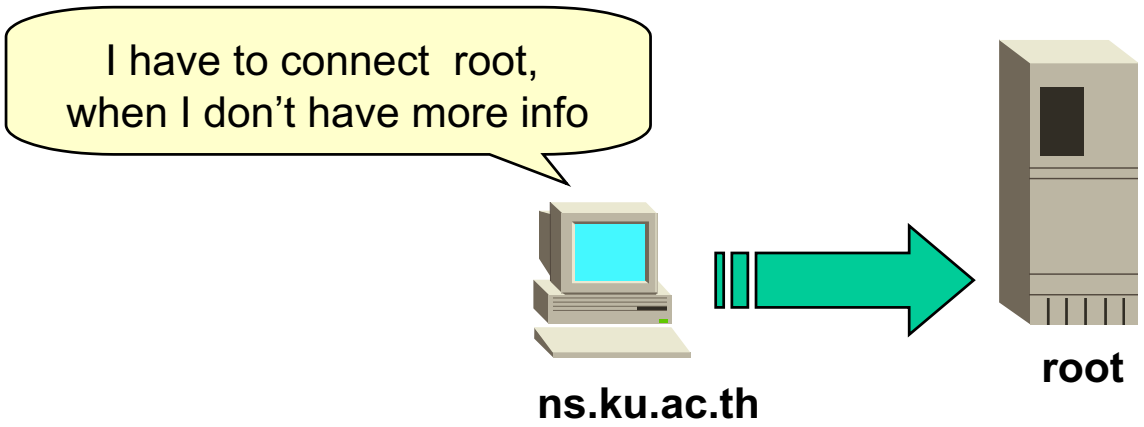

# Root Name Server (cont.)

- 13 root servers are currently available in Internet (Last updated Aug 97)

| | |
|---|---|
| A.ROOT-SERVERS.NET | 198.41.0.4 |
| B.ROOT-SERVERS.NET | 128.9.0.107 |
| C.ROOT-SERVERS.NET | 192.33.4.12 |
| D.ROOT-SERVERS.NET | 128.8.10.90 |
| E.ROOT-SERVERS.NET | 192.203.230.10 |
| F.ROOT-SERVERS.NET | 192.5.5.241 |
| G.ROOT-SERVERS.NET | 192.112.36.4 |
| H.ROOT-SERVERS.NET | 128.63.2.53 |
| I.ROOT-SERVERS.NET | 192.36.148.17 |
| J.ROOT-SERVERS.NET | 198.41.0.10 |
| K.ROOT-SERVERS.NET | 193.0.14.129 |
| L.ROOT-SERVERS.NET | 198.32.64.12 |
| M.ROOT-SERVERS.NET | 202.12.27.33 |

ftp://ftp.rs.internic.net/domain/named.root

# Name Resolution Process

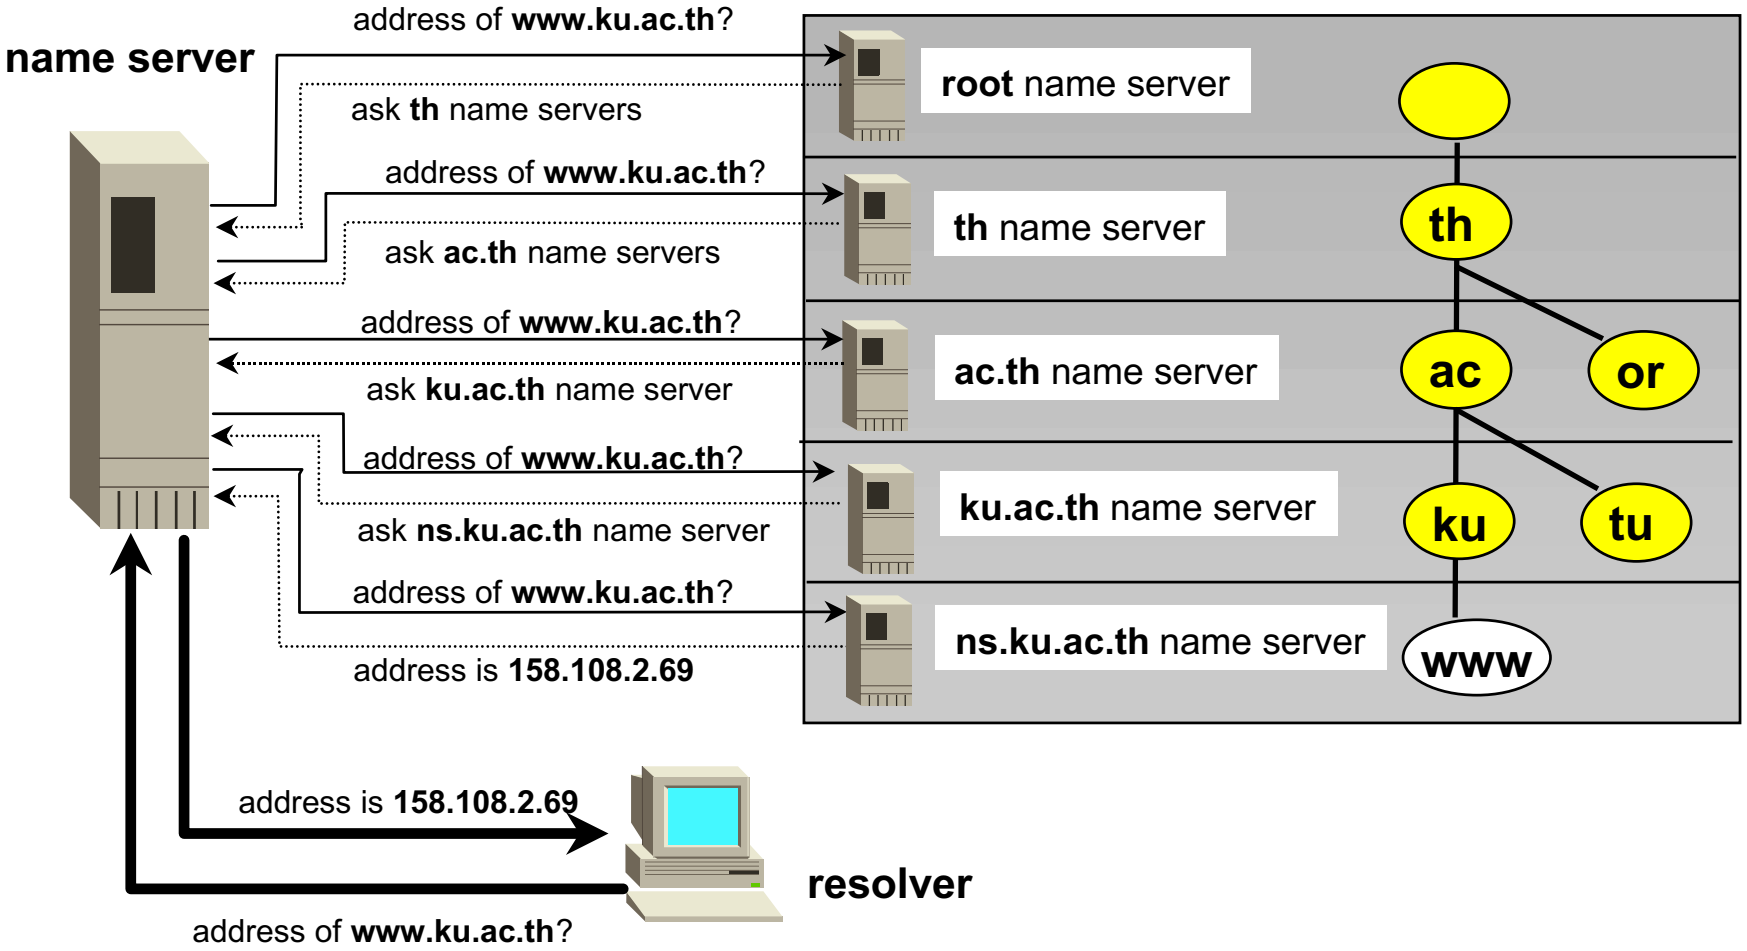

# Reverse Resolution

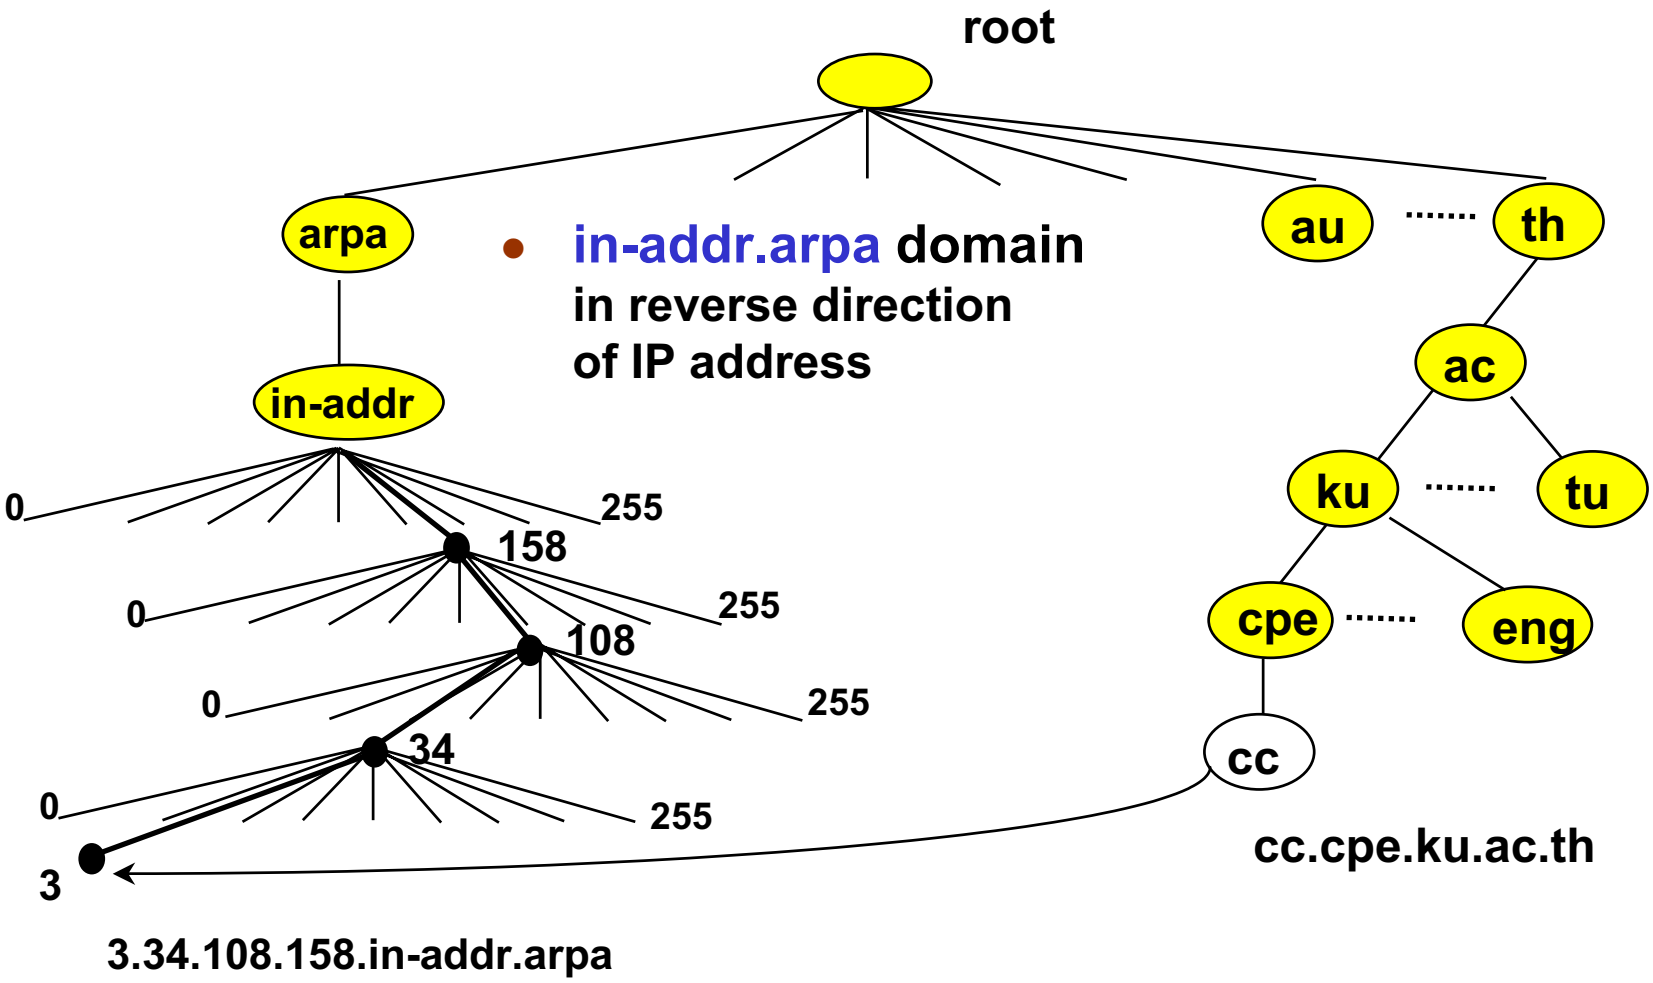


- **in-addr.arpa domain** in reverse direction of IP address

# Caching

- all name servers employ a cache to reduce the DNS traffic
- standard UNIX keep cache in name server with time-out
- cache data is *non-authoritative*

# DNS message format (I)

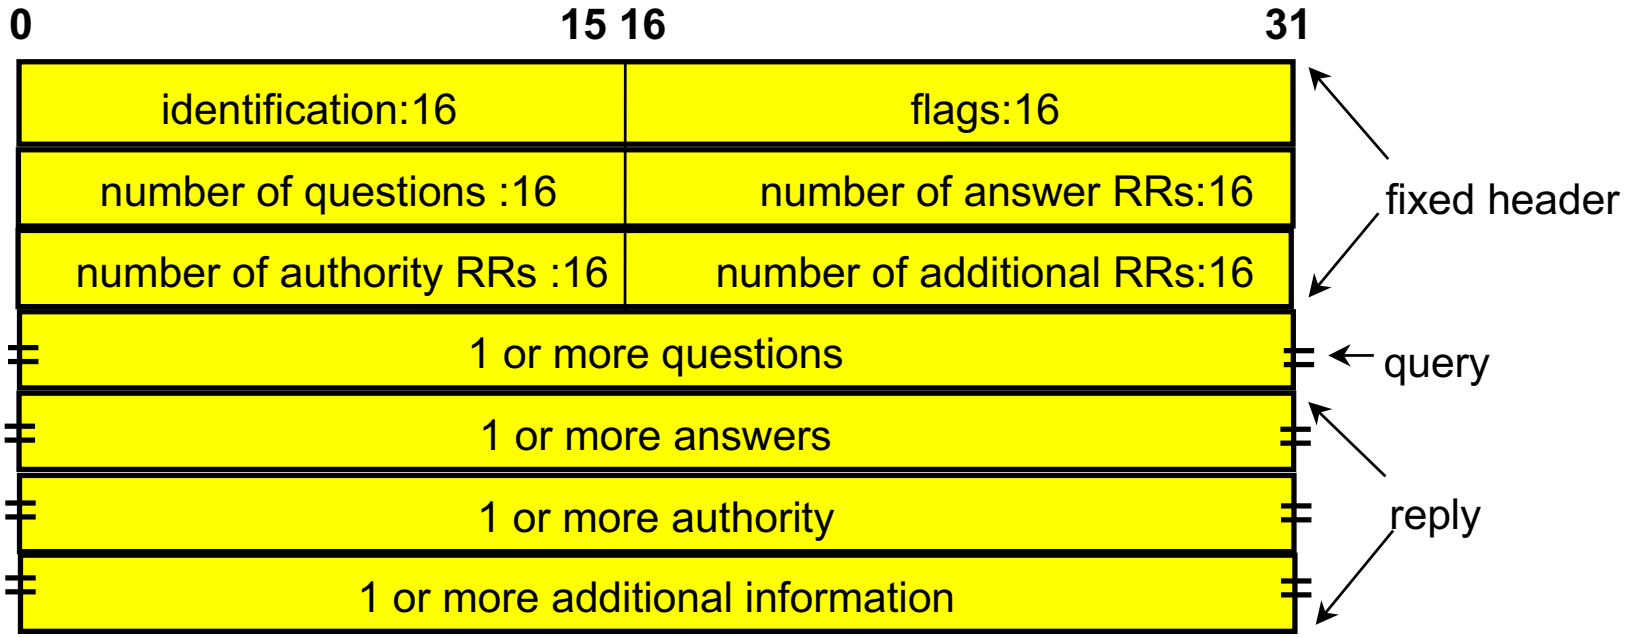


- fixed 12 byte header with 4 variable length fields
- DNS message format is defined for both queries and answers

# DNS message format (II)

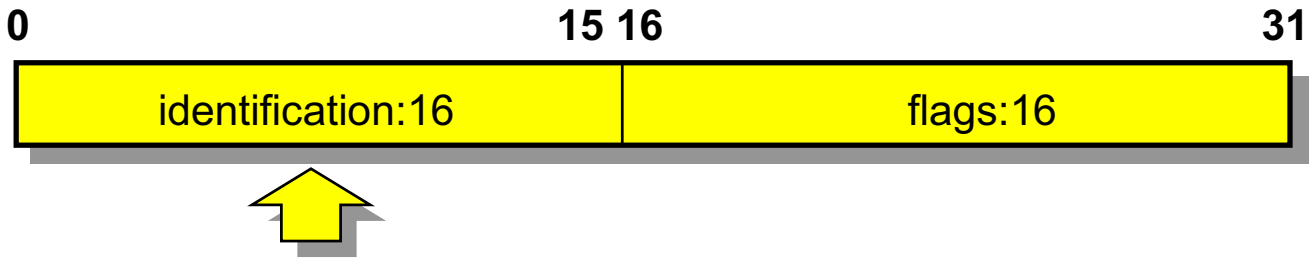


- set by the client and return by the server
- lets the client match responses to requests

# DNS message format (III)

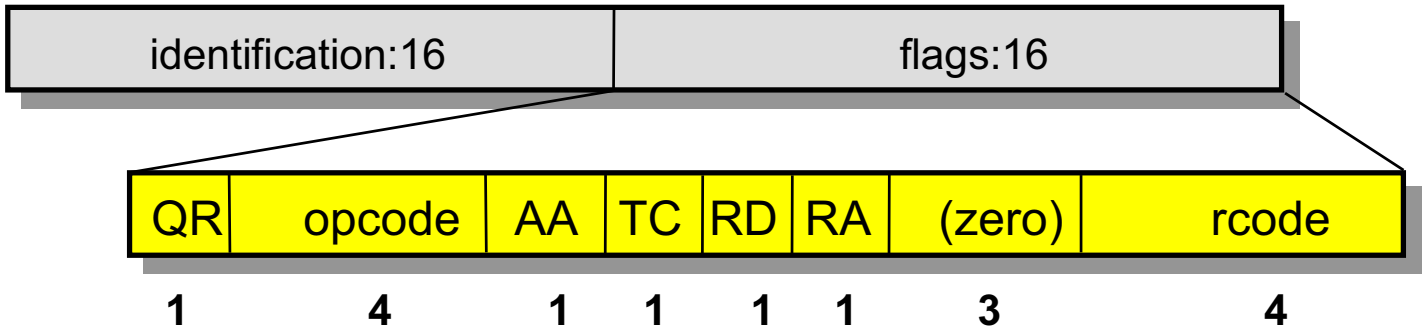


| | |
|---|---|
| QR | 0= query, 1= response |
| opcode | 0= standard query, 1=inverse query, 2=server status request |
| AA | 0= authoritatived answer,1 = non authoritatived answer |
| TC | 1= truncated. using UDP, reply was>512 bytes, return only 512 bytes |
| RD | 1= recursive desired, 0= iterative |
| RA | 1= recursion available (server support recursion) |
| rcode | return code : 0=no error, 3=name error |

# DNS message format (IV)

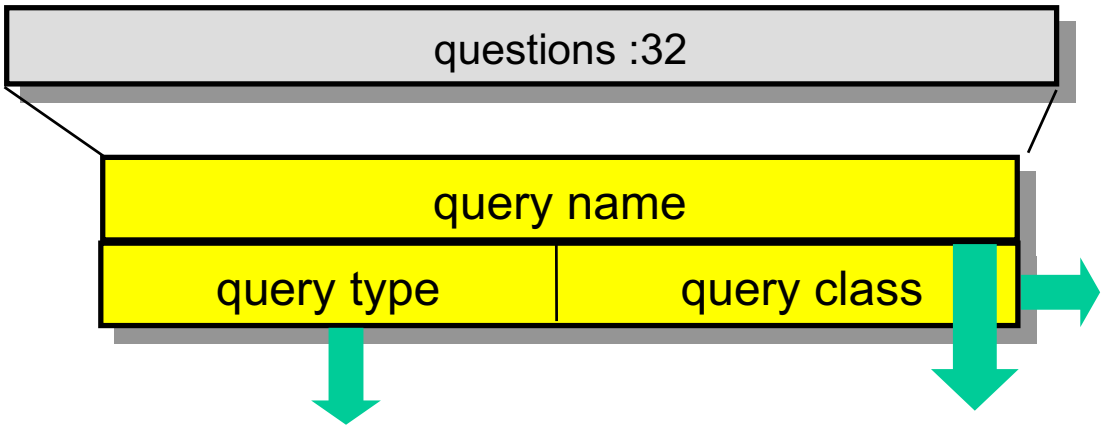


***query class*** **normally =1 means Internet Address**

***query type*** **indicates desired response**

| | | |
|---|---|---|
| A | 1 | IP address |
| NS | 2 | name server |
| CNAME | 5 | canonical name |
| PTR | 12 | pointer record |
| HINFO | 13 | host info |
| MX | 15 | mail exchange record |

***query name*** **is the name being lookup, sequence of label begins with 1-byte count**

**3www3cpe2ku2ac2th0**

# DNS message format (V)

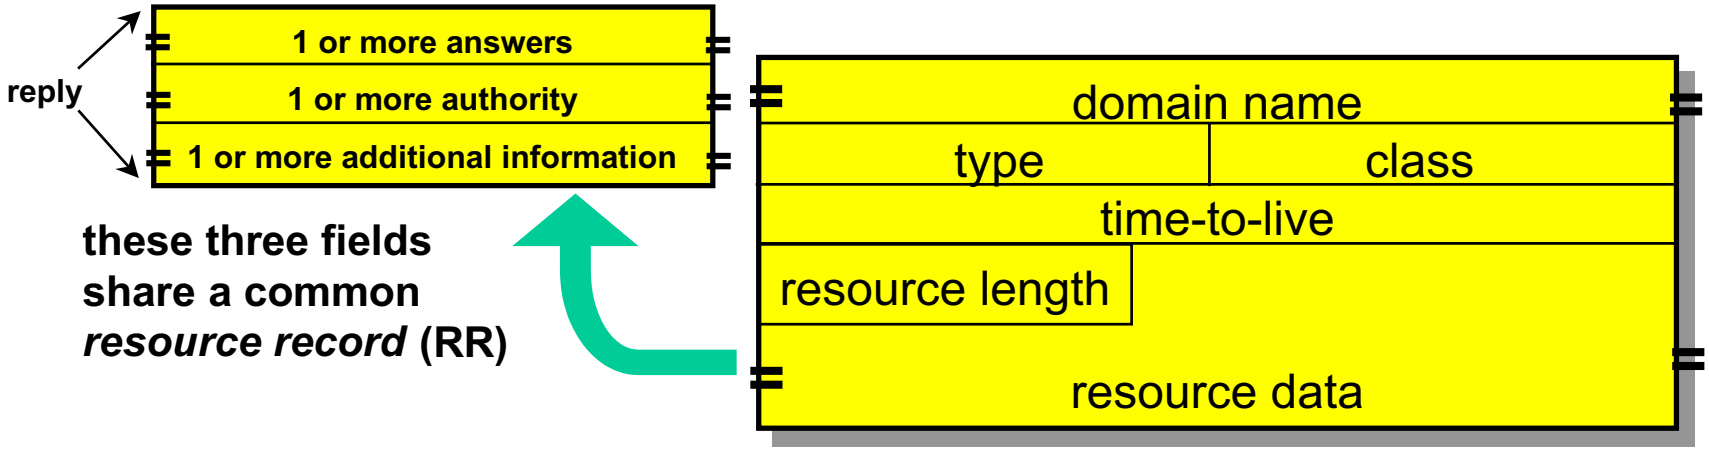


- *domain name* : corresp. response name, (*query name* format)
- *type* : response RR type code (see *query type*)
- *time-to-live* : cache life time of RR (often 86400=1 day)
- *resource length* : specify the size of resource data
- *resource data* : the answer, e.g. IP address or other type

# Operations

- use port 53
- typically UDP request and reply
- if answer is too big, use TCP

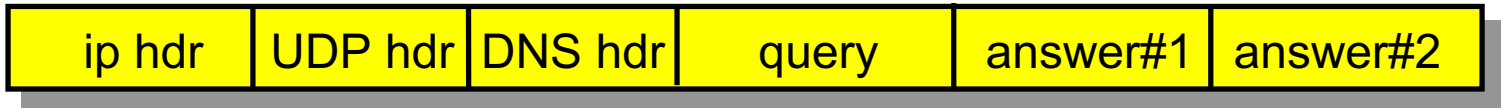

# Resolver file

- **resolver must have address for local name server**
- ***/etc/resolv.conf* on UNIX**

***/etc/resolv.conf***

```
# domain
domain  cpe.ku.ac.th
# list of name server
nameserver 158.108.32.5
nameserver 158.108.33.3
nameserver 158.108.2.67
```

# Setting up DNS

- **BIND (Berkeley Internet Name Domain) package**
- ***/usr/somewhere/in.named* - BSD named DNS server**
- ***/etc/named.boot* - named configuration (tell named where to find database files)**

# Sample named.boot

```
; Boot file for server ns.cpe.ku.ac.th.
directory      /usr/local/named
cache          .                        root.cache
primary        localhost.               primary/local
primary        0.0.127.in-addr.arpa     primary/local.rev
primary        cpe.ku.ac.th             primary/cpe
secondary      ku.ac.th                 158.108.2.67      secondary/ku
secondary      cpc.ku.ac.th             158.108.2.67      secondary/cpc
secondary      eng.ku.ac.th             158.108.40.196    secondary/eng
secondary      108.158.in-addr.arpa     158.108.2.67      secondary/ku.rev
primary        32.108.158.in-addr.arpa  primary/zone/zone32
primary        33.108.158.in-addr.arpa  primary/zone/zone33
   :                    :                        :
```

# New named.conf format

- **BIND Version 8 defines a new format of boot file : *named.conf***

```
options {
        directory "/usr/local/named";
};

zone "." {
        type hint;
        file "root.cache";
};

zone "localhost." {
        type master;
        file "primary/local";
};
```

```
zone "0.0.127.in-addr.arpa" {
        type master;
        file "primary/local.rev";
};

zone "cpe.ku.ac.th" {
        type master;
        file "primary/cpe";
};

zone "ku.ac.th" {
        type slave;
        masters {  158.108.2.67;
        };
:
:
```

# Sample database file

```
; address file for server ns.cpe.ku.ac.th. (primary/cpe)
@           IN    SOA    ns.cpe.ku.ac.th. dnsadmin.ns.cpe.ku.ac.th. (
                         1998051300     ; Last Updated May 13,1998
                         10800          ; Refresh every 3 hours
                         3600           ; Retry every 1 hour
                         2592000        ; Expire after 30 days
                         86400          ; Minimum TTL of 1 day
                         )
; Name Servers
            IN    NS     ns.cpe.ku.ac.th.
            IN    NS     ns.eng.ku.ac.th.
            IN    NS     ns.ku.ac.th.
```

(*continue on next page*)

# Sample database file (cont.)

```
; Mail Hubs for the Domain
            IN      MX      10      mailhost.cpe.ku.ac.th.
            IN      MX      15      cc.cpe.ku.ac.th.
; AI Lab
mars        IN      HINFO       "MP 1101D/DECstation 5000-420"
            IN      MX          10      mailhost.cpe.ku.ac.th.
            IN      A           158.108.32.97
maspar      IN      CNAME       mars.cpe.ku.ac.th.

saturn      IN      HINFO       "SPARCstation 2" "SunOS 4.1.3"
            IN      MX          10 mailhost.cpe.ku.ac.th.
            IN      A           158.108.32.98
ailab2      IN      CNAME       saturn.cpe.ku.ac.th.
```