Internet Engineering

241-461

Robert Elz

kre@munnari.OZ.AU
kre@coe.psu.ac.th

http://fivedots.coe.psu.ac.th/~kre

Contents

- The Domain Name System
- The DNS Database
- DNS Protocols
- DNS Message Formats
  - Queries
  - Responses
- DNS Limits
- Zone Transfer
- Mapping Addresses to Names

DNS Message Format

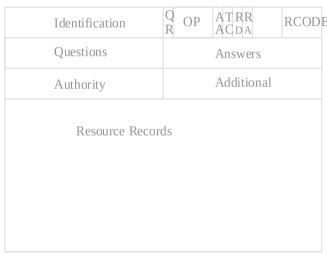

- Header (12 bytes)
- Some number of variable length RRs

DNS Message Sections

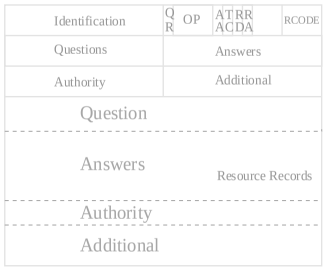

- Header fields give RR counts in each section
- RRs not necessarily multiples of 4 bytes
  - nor are sections

RR Format

- Label: The domain name
- Class and Type - classify the particular data
- TTL: how long from when it was received the data is valid

Except in the question section:

RR Label

- Sequence of explicitly counted strings
- Zero count means the end of the sequence
  - Hence no intermediate label
  - can be shorter than 1 byte
- Labels are any binary data
  - except ASCII case is ignored
  - A == a B == b ...

RR Label Length

- Length in a 6 bit field
  - Max label length is 63 bytes
- Values other than 0 0 in top two bits
  - indicate something different
- Total length limited to 255 bytes
  - not required by message formats

DNS Queries

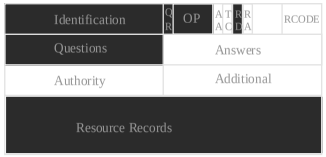

- QR bit: Query or Response
  - 0 => Query, 1 => Response
- Op: Opcode
  - 0 for all normal DNS queries/replies
- RD bit: Recursion Desired
  - Set implies
    - we want the server to find the answer
  - Reset implies
    - we just want server to tell us what it knows

DNS Queries (2)

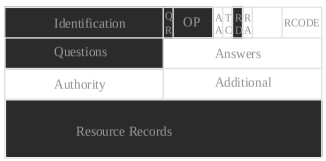

- Identification:
  - a number picked by the resolver
  - allows the reply
    - to be easily associated with the query
- Questions:
  - count of RR's in question section
    - almost always 1
- RR's:
  - Question section,
    - the (abbreviated) RR specifying the
      - name, class, type,
    - for which information is sought

DNS Replies

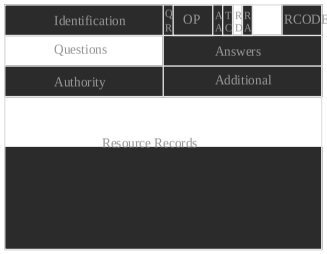

- QR bit: set to 1 in a reply.
- Many fields copied from Query:
  - Op, Identification, Questions, RD bit

DNS Replies (2)

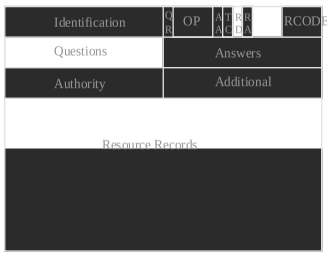

- rcode: reply code (error indicator)
  - 0 ==> no error
  - other values various errors

DNS Replies (3)

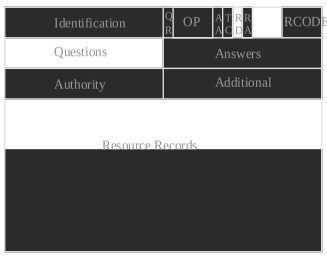

- AA: Authoritative Answer
  - Server Authoritative for Zone in Question
  - Answer was not extracted from cached data
- RA: Recursion Available
  - Server is willing to act as back-end resolver
- TC: Truncated Reply
  - Answer did not fit in packet

DNS Replies (4)

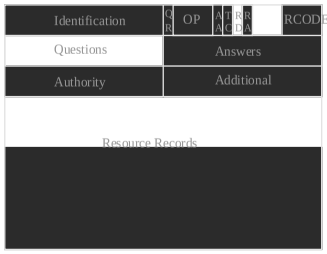

- Answers:
  - Specific answers to the questions asked

DNS Replies (5)

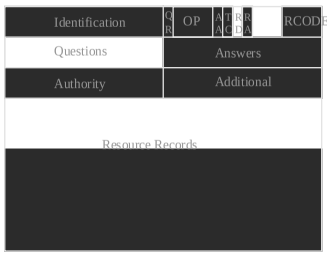

- Authority:
  - Indication where the authoritative servers are
  - For negative answers, who says so.
    - That is, for errors

DNS Replies (6)

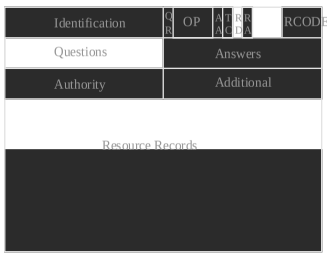

- Additional:
  - Other information that might be useful
    - Addresses of names mentioned (etc).

DNS Examples

Query to find address of fivedots.coe.psu.ac.th

```
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 4
;; flags: qr aa rd ra;
;; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 2
;; QUERY SECTION:
;;      fivedots.coe.psu.ac.th, type = A, class = IN

;; ANSWER SECTION:
fivedots.coe.psu.ac.th.  1W IN A  172.30.0.5

;; AUTHORITY SECTION:
coe.psu.ac.th.          1W IN NS        ninedots.coe.psu.ac.th.
coe.psu.ac.th.          1W IN NS        fivedots.coe.psu.ac.th.

;; ADDITIONAL SECTION:
ninedots.coe.psu.ac.th.  1W IN A  172.30.0.9
fivedots.coe.psu.ac.th.  1W IN A  172.30.0.5
```

DNS Examples (2)

Query to find address of www.psu.ac.th

```
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 4
;; flags: qr aa rd ra;
;; QUERY: 1, ANSWER: 2, AUTHORITY: 2, ADDITIONAL: 2
;; QUERY SECTION:
;;      www.psu.ac.th, type = A, class = IN

;; ANSWER SECTION:
www.psu.ac.th.          12H IN CNAME    s1.psu.ac.th.
s1.psu.ac.th.           12H IN A        192.100.77.61

;; AUTHORITY SECTION:
psu.ac.th.              12H IN NS       ns.psu.ac.th.
psu.ac.th.              12H IN NS       ns2.psu.ac.th.

;; ADDITIONAL SECTION:
ns.psu.ac.th.           12H IN A        192.100.77.5
ns2.psu.ac.th.          12H IN A        192.100.77.2
```

DNS Examples (3)

Query to find address of oddball.psu.ac.th

```
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 4
;; flags: qr aa rd ra;
;; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 0
;; QUERY SECTION:
;;      oddball.psu.ac.th, type = A, class = IN

;; AUTHORITY SECTION:
psu.ac.th.  12H IN SOA   (
        ns.psu.ac.th. netadmin.ratree.psu.ac.th.
        2001080801      ; serial
        10800           ; refresh -  3 hours
        1200            ; retry   - 20 minutes
        1209600         ; expiry  -  2 weeks
        43200 )         ; minimum - 12 hours
```

DNS Query Types

- Can Query for Specific RR Type
  - A - IPv4 Address
  - MX - Mail Exchange
  - NS - Name Servers
  - CNAME - Canonical Name
    - Lots more
- Can Query for Meta-Type
  - ANY - Any RR's owned by the name
  - AXFR - All RR's in Zone Named

DNS Responses

- No Error
  - Answer Provided
    - ...
  - No Answer Provided
    - No data of requested type exists
  - Referral
    - Try again at one of these servers
- Error
  - Name Invalid
  - Server Error
  - Format Error
  - Not Implemented
  - Refused