CS 43: Computer Networks

Writing a DNS Client

9 October, 2020

Courtesy: Kurose & Ross, K. Webb, Stanford University

DNS Message Structure

Has the same format for query and response

Header
Question
Answer
Authority
Additional

Query only has the
Header and Question portions

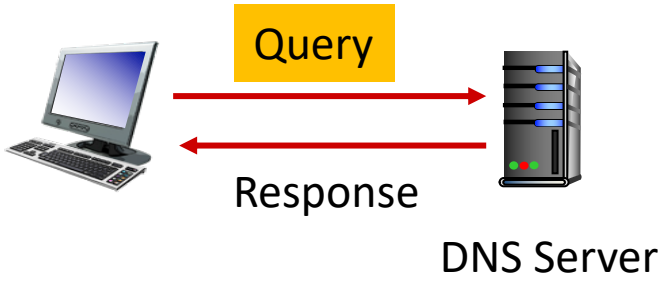

DNS Message Structure

Has the same format for query and response

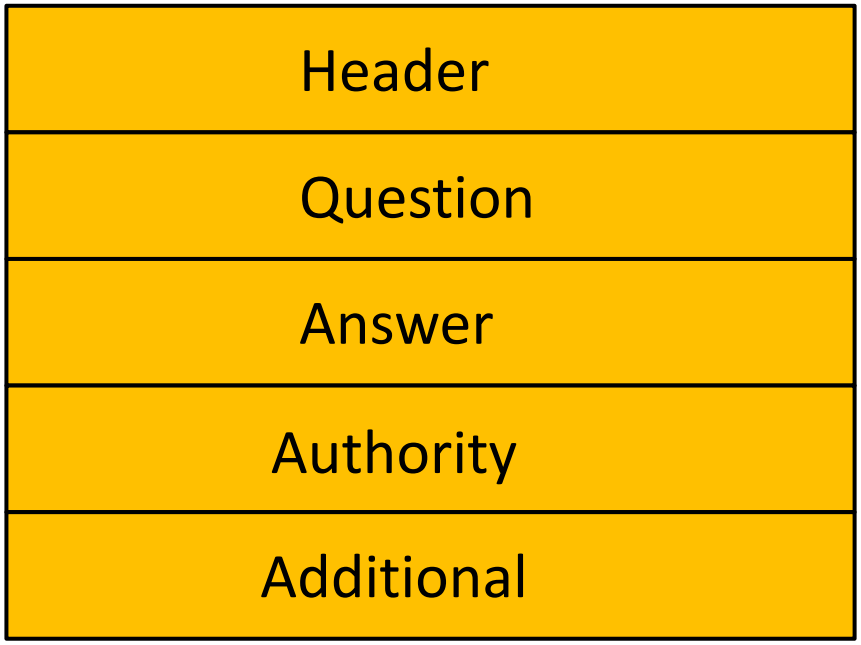

Query only has the
Header and Question portions

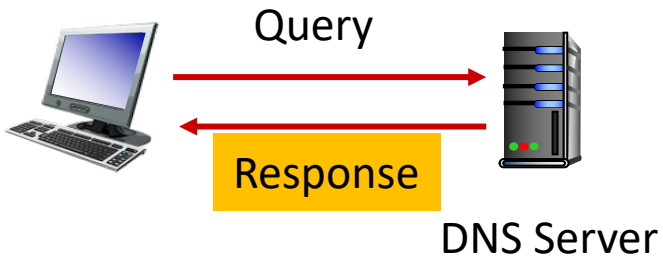

DNS Message Header

Has the same format for query and response

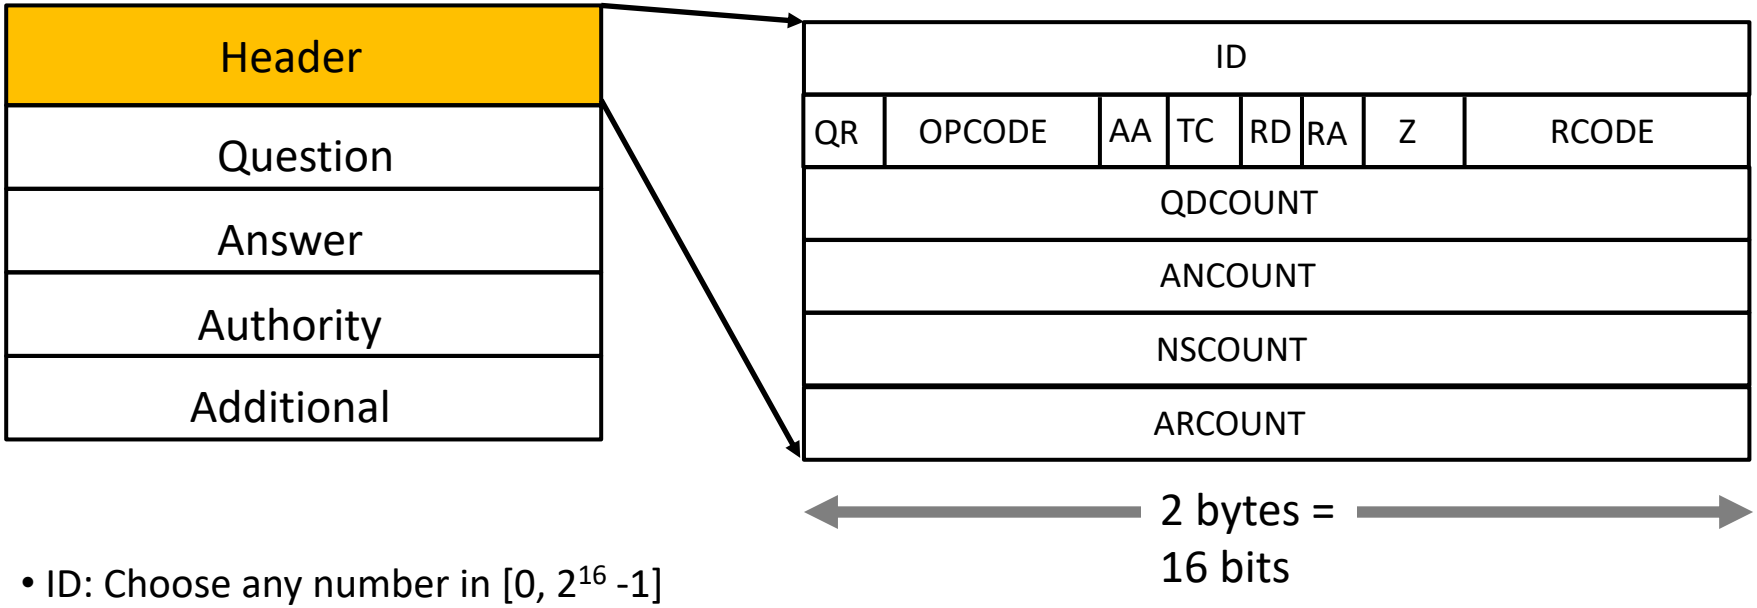

- ID: Choose any number in [0, 2^{16} -1]
- QR: Query = 0, Response = 1
- OPCODE: Standard query = 0
- RCODE: Error Code
- Flags:
  - AA: authoritative server
  - TC: Truncated
  - RD: Recursion Desired
  - RA: Recursion Available

- QDCOUNT: Number of Questions
- ANCOUNT: Number of Answers
- NSCOUNT: Number of Name Server Records
- ARCOUNT: Number of Additional Resource Records

DNS Message Header

Example query dig demo.cs.swarthmore.edu

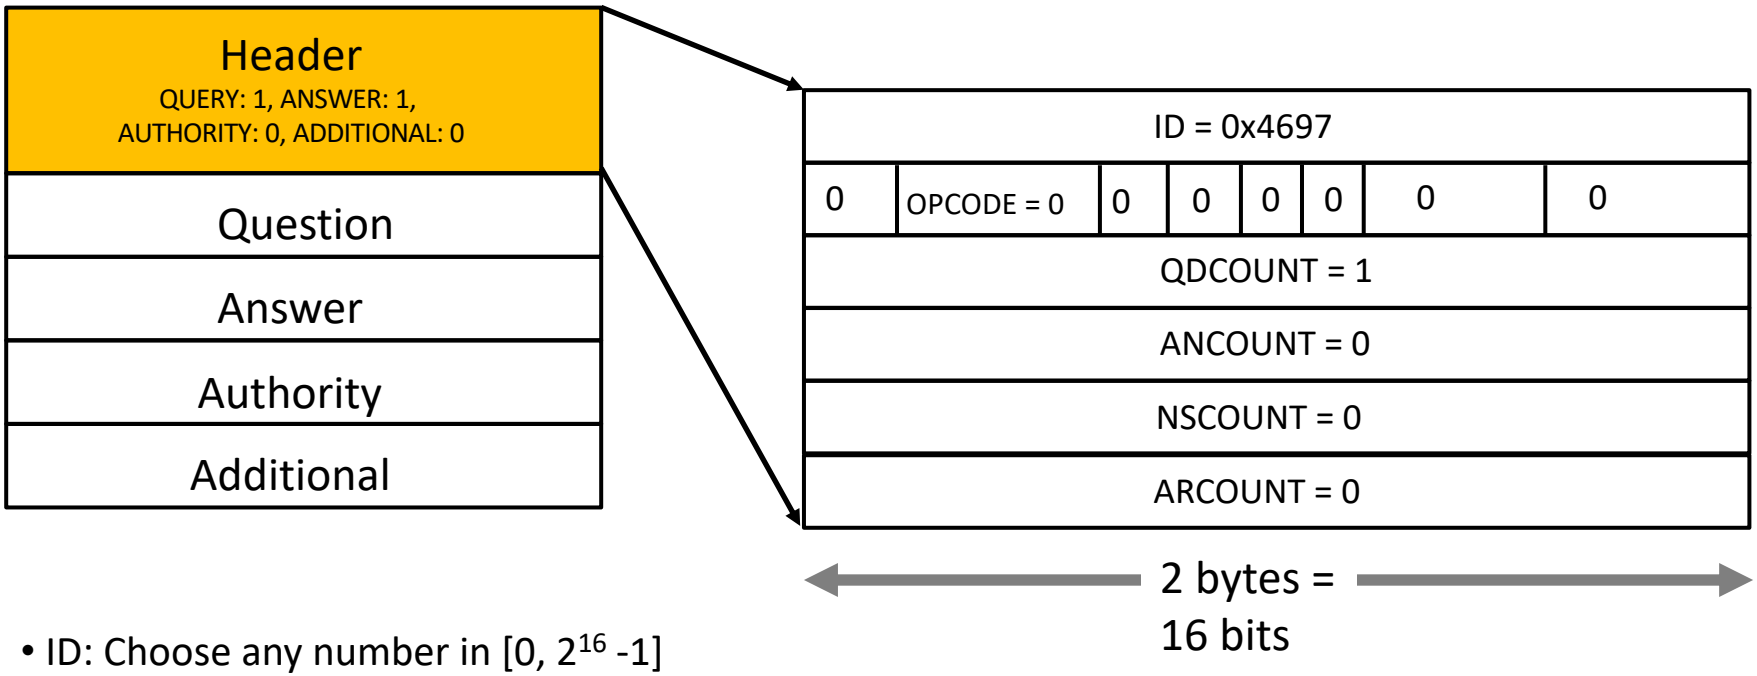

- ID: Choose any number in [0, 2^{16} -1]
- QR: Query = 0, Response = 1
- OPCODE: Standard query = 0
- RCODE: Error Code
- Flags:
  - AA: authoritative server
  - TC: Truncated
  - RD: Recursion Desired
  - RA: Recursion Available

- QDCOUNT: Number of Questions
- ANCOUNT: Number of Answers
- NSCOUNT: Number of Name Server Records
- ARCOUNT: Number of Additional Resource Records

DNS Query format

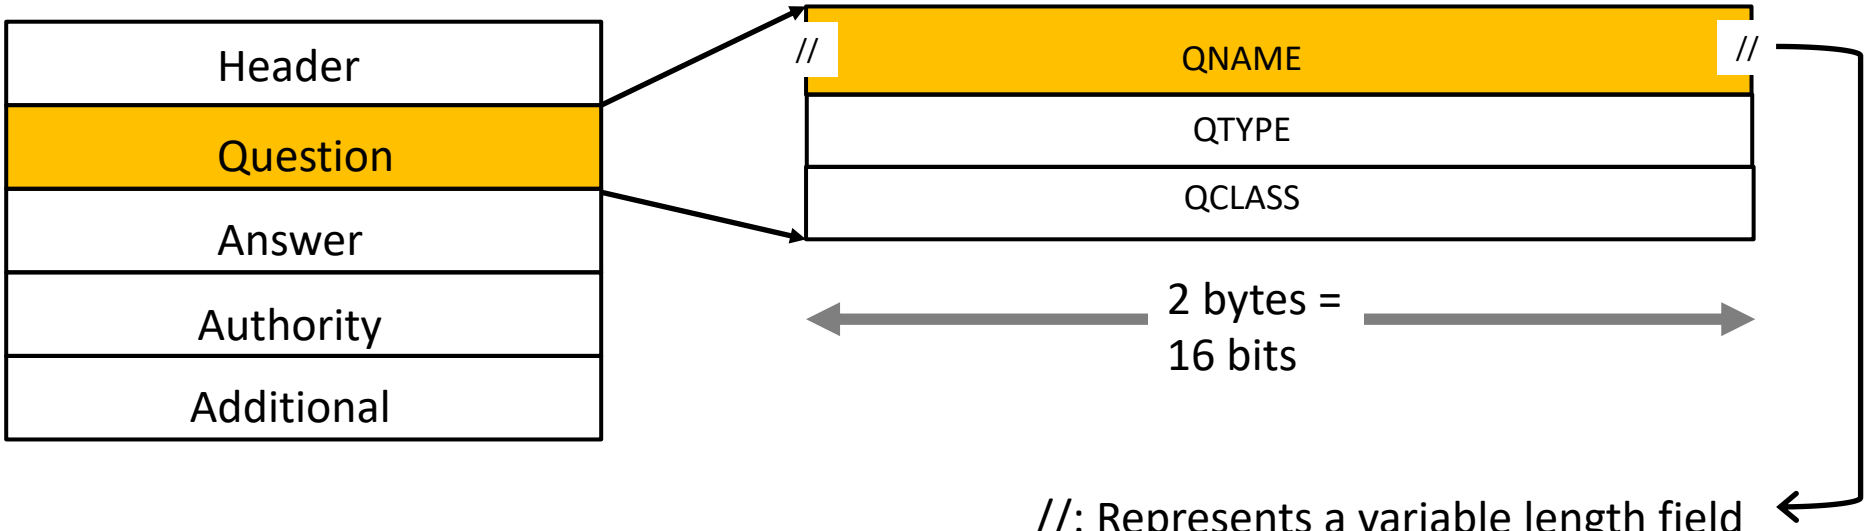

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 , MX = 15..
- CLASS: Internet = 1

//: Represents a variable length field
QNAME= variable length
QTYPE = 2 bytes
QCLASS = 2 bytes

DNS Message Resource Record format

All of the resource records follow the same format

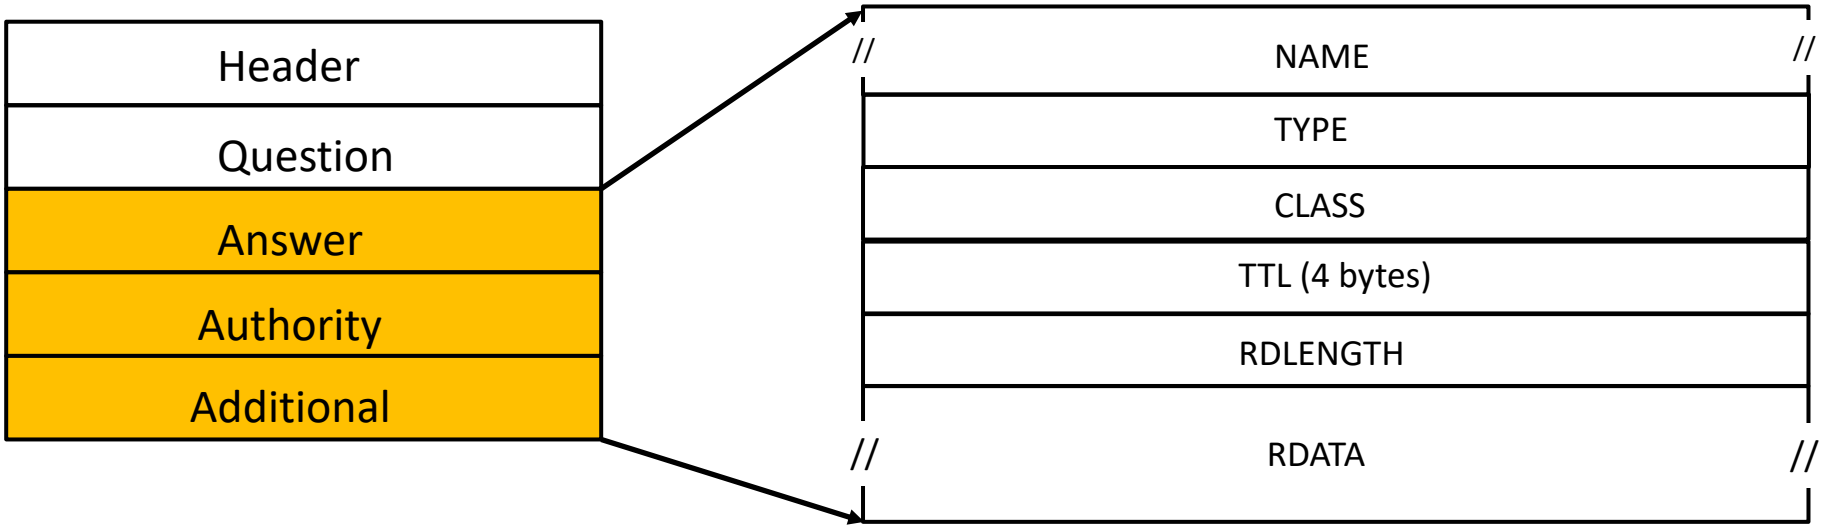

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 ..
- CLASS: Internet = 1
- TTL: time-to-live (seconds)
- RDLENGTH: length of RDATA
- RDATA: record data (variable length)

DNS response format

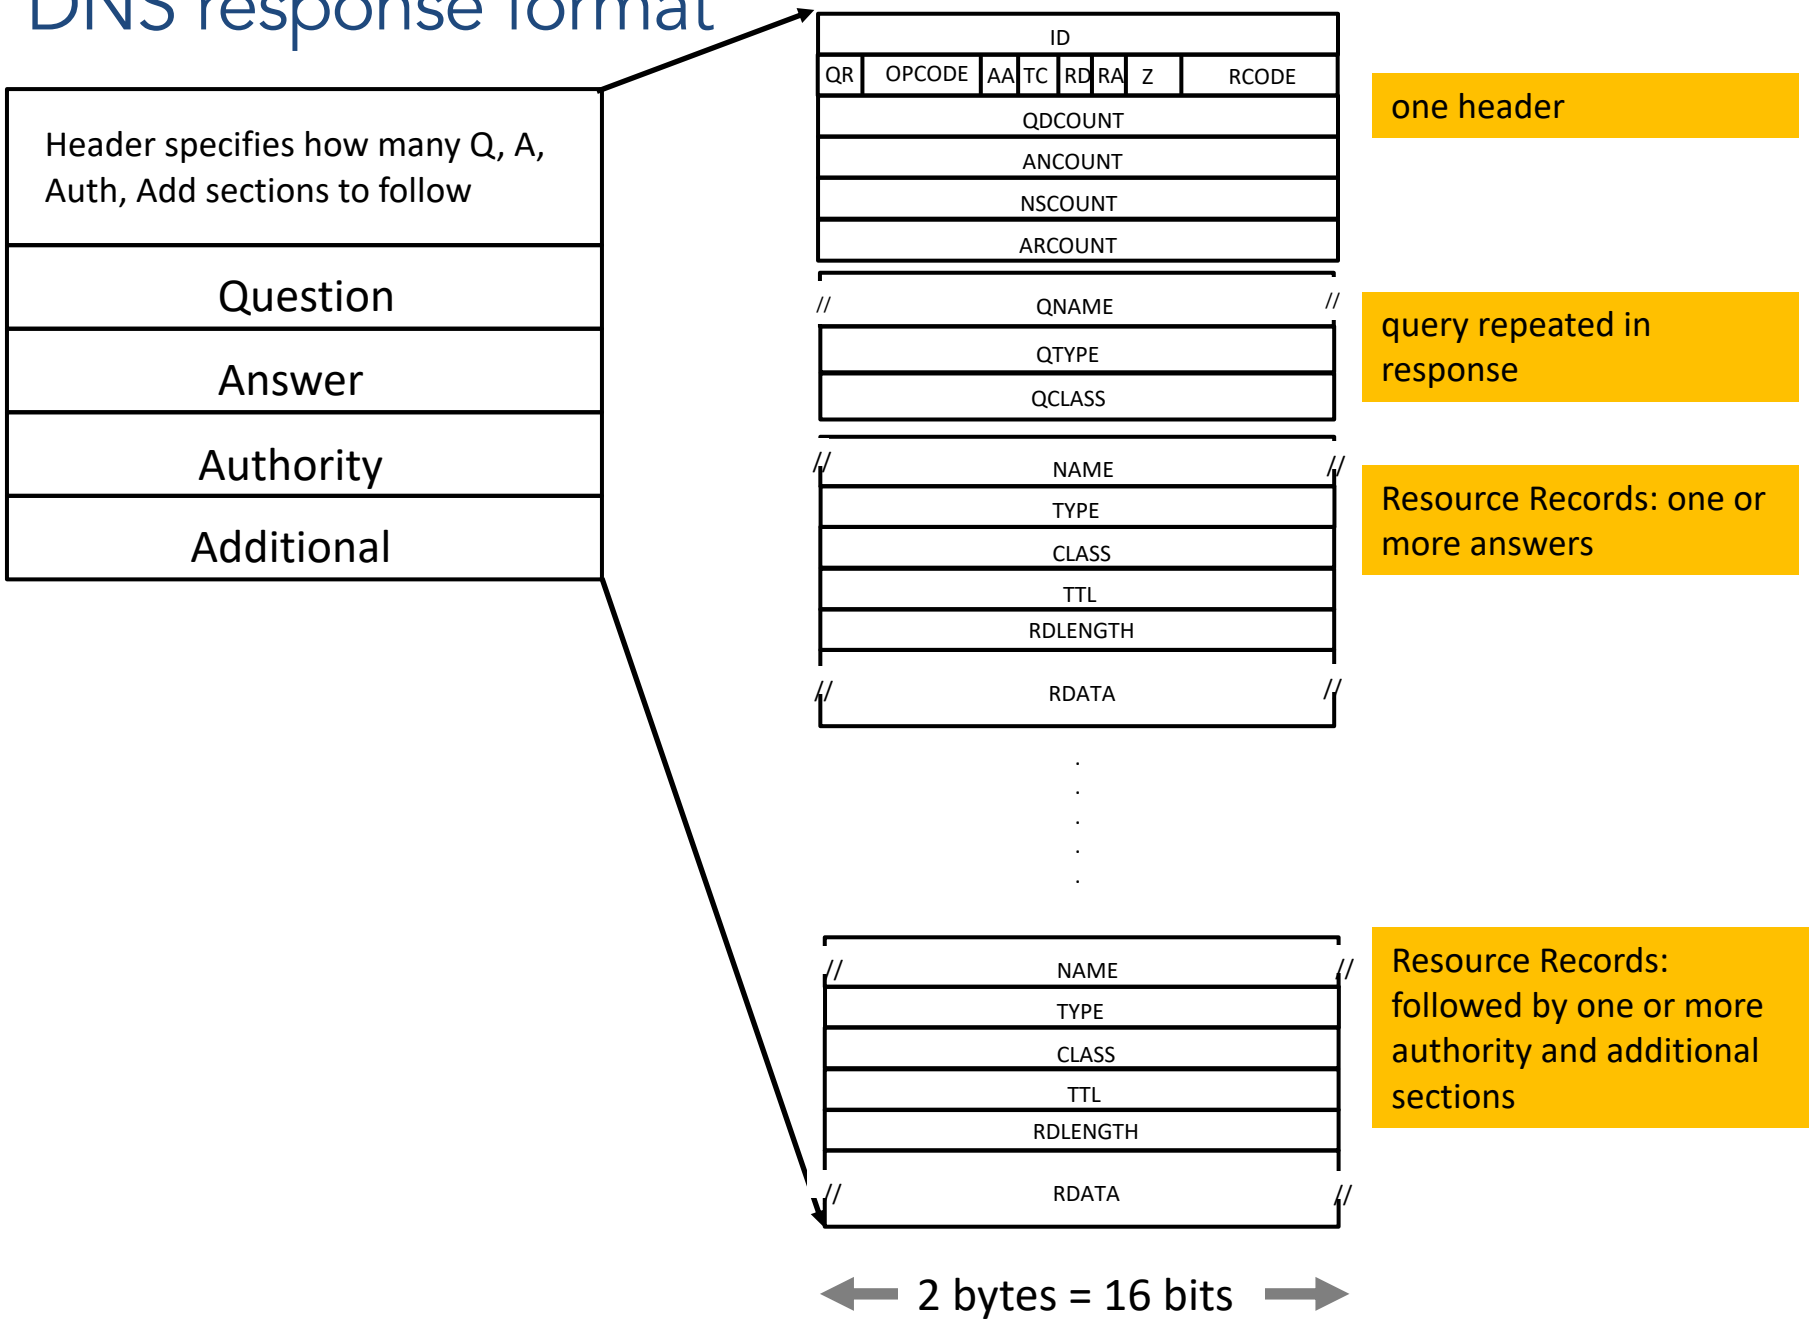

DNS Query format

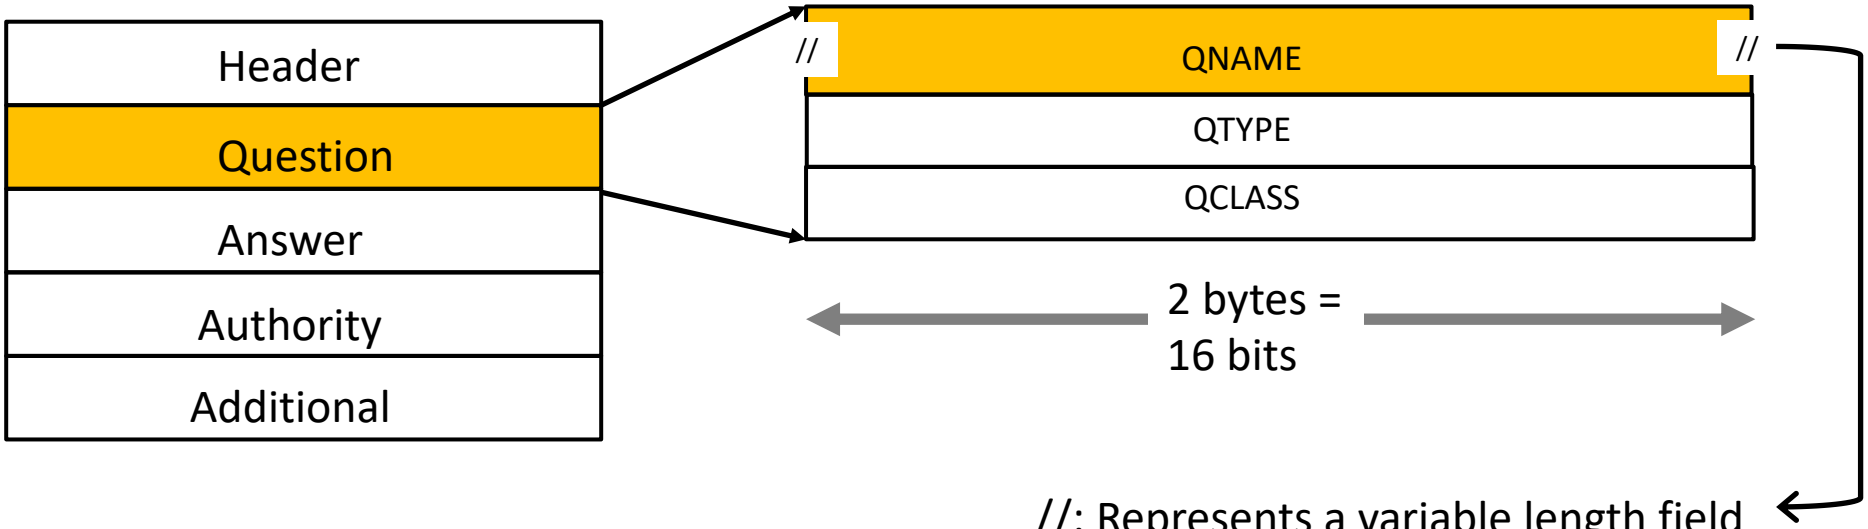

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 , MX = 15..
- CLASS: Internet = 1

//: Represents a variable length field
QNAME= variable length
QTYPE = 2 bytes
QCLASS = 2 bytes

DNS Name Encoding

- Names can be long and repeated several times in a packet.
  - Query/Answer
  - NS record/ A record
- Name semantics: Break name into labels
- demo.cs.swarthmore.edu. => 4 demo 2 cs 10 swarthmore 3 edu 0 (nameless root)
  - length fields are in binary, text in ASCII
  - e.g. 4 demo in hex = 04 64 65 6d 6f
  - periods in between the hostname (.) are not included

DNS Query format

Example query:

dig +norec demo.cs.swarthmore.edu @ibext.its.swarthmore.edu

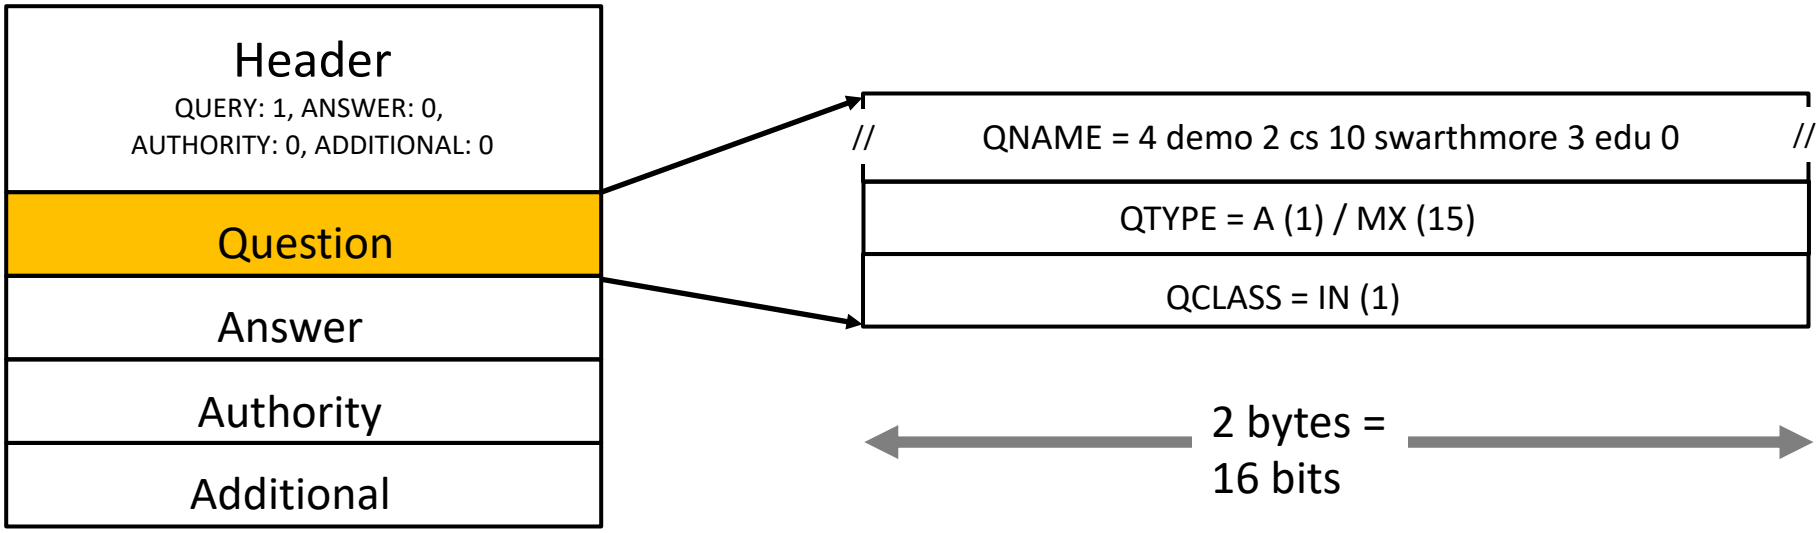

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 , MX = 15..
- CLASS: Internet = 1

//: Represents a variable length field
QNAME= variable length
QTYPE = 2 bytes
QCLASS = 2 bytes

DNS A Resource Record
Hostname to IP address mapping

dig +norec demo.cs.swarthmore.edu @ibext.its.swarthmore.edu

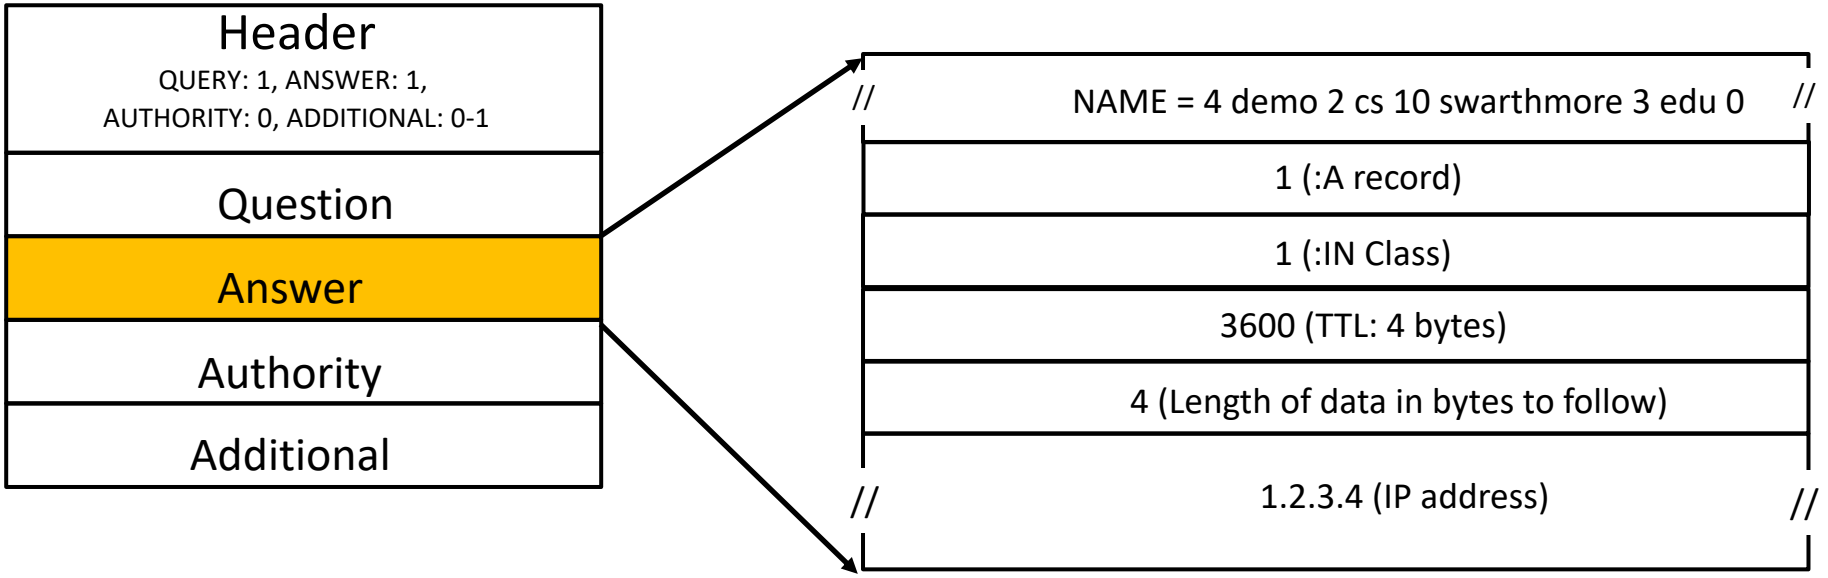

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 ..
- CLASS: Internet = 1
- TTL: time-to-live (seconds)
- RDLENGTH: length of RDATA
- RDATA: record data (variable length)

DNS NS Resource Record
Hostname to DNS server mapping

dig +norec demo.cs.swarthmore.edu @a.root-servers.net

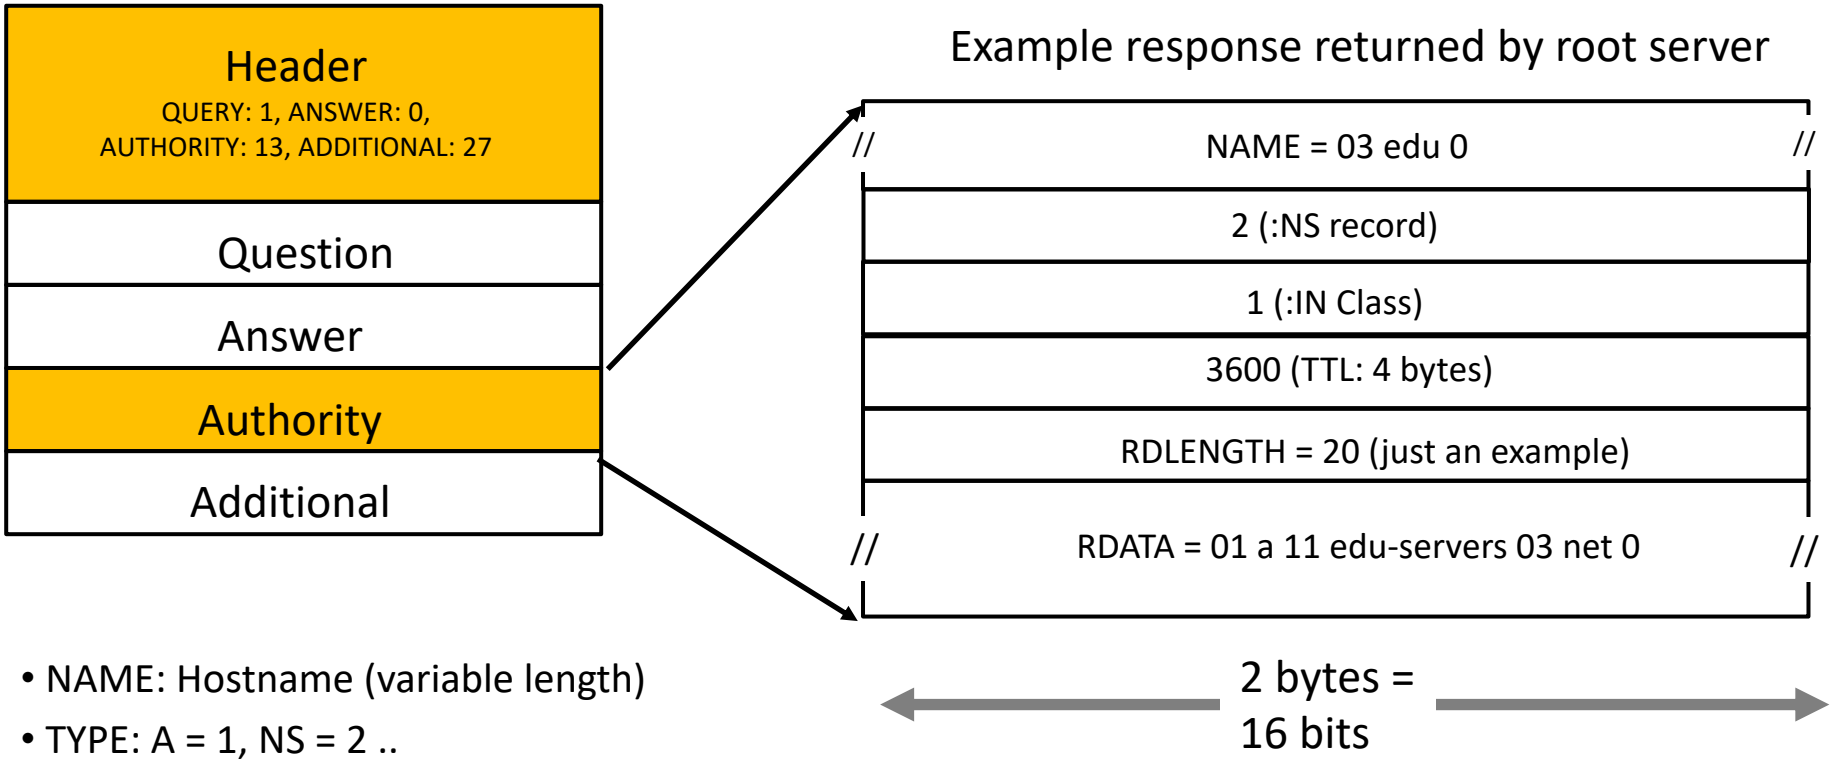

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 ..
- CLASS: Internet = 1
- TTL: time-to-live (seconds)
- RDLENGTH: length of RDATA
- RDATA: record data (variable length)

DNS NS Resource Record
Hostname to DNS server mapping

dig +norec demo.cs.swarthmore.edu @a.root-servers.net

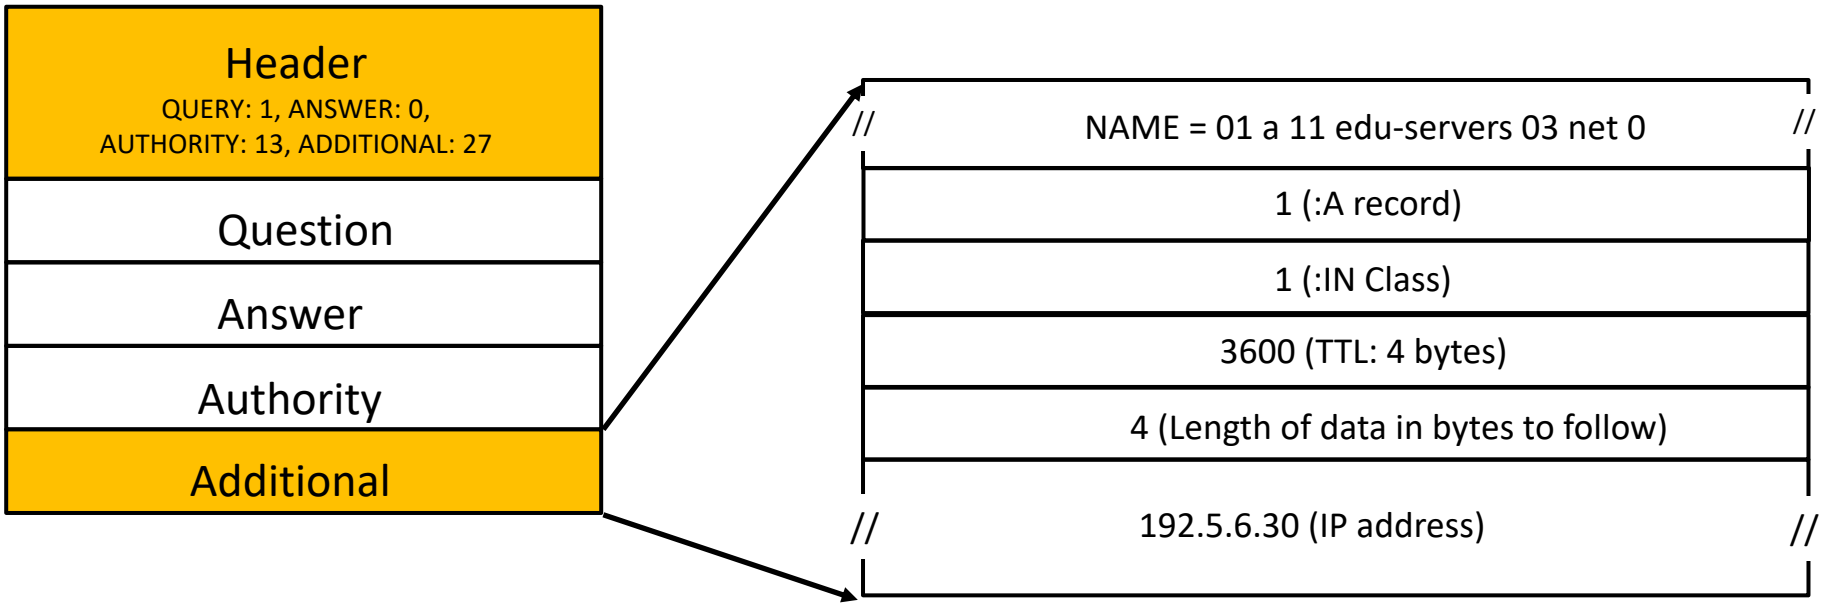

- NAME: Hostname (variable length)
- TYPE: A = 1, NS = 2 ..
- CLASS: Internet = 1
- TTL: time-to-live (seconds)
- RDLENGTH: length of RDATA
- RDATA: record data (variable length)

DNS response format

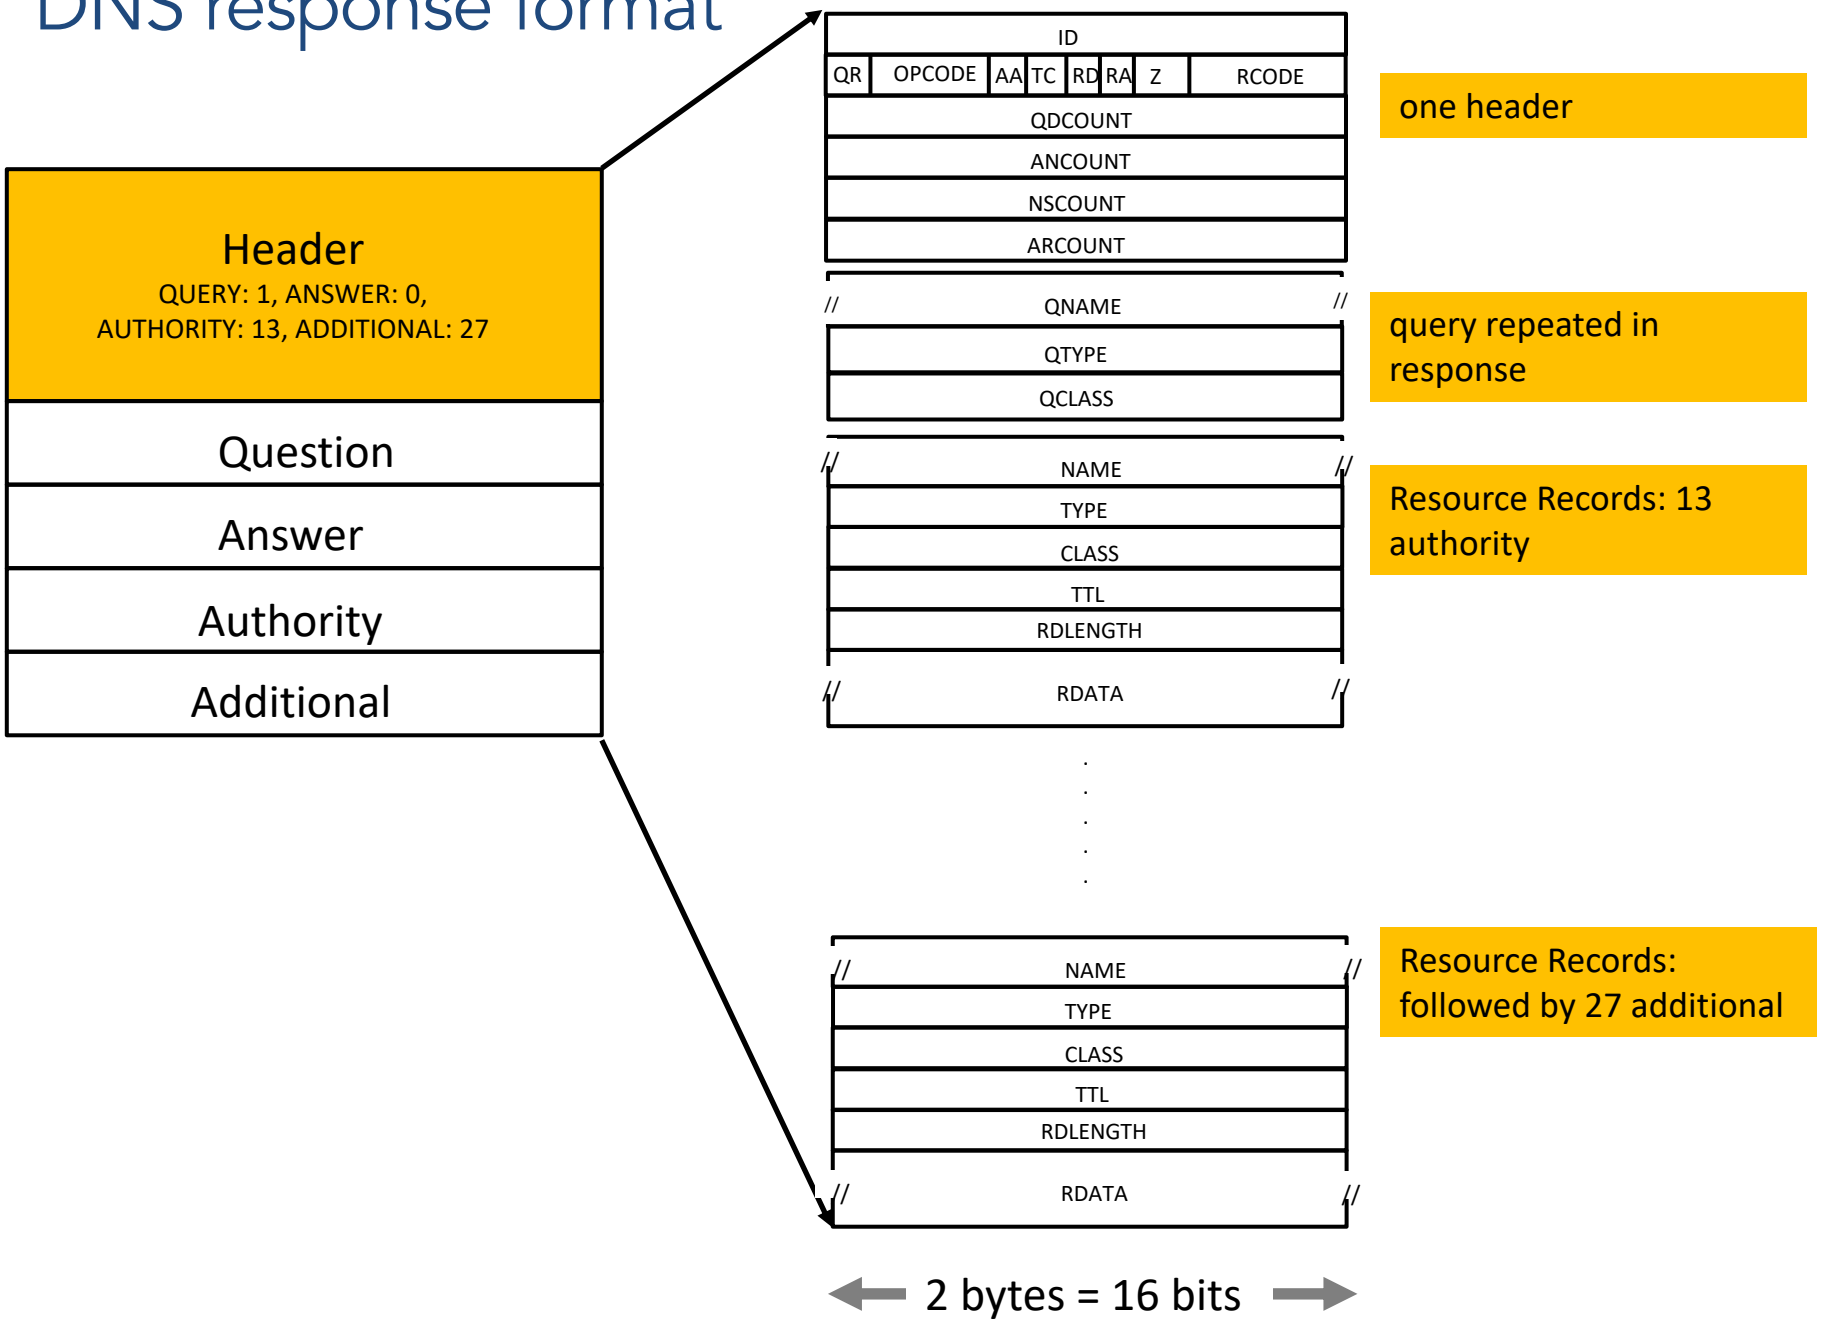

DNS Name Compression

slightly more complicated..

Since names can several times in a packet, DNS uses name compression in for text fields.

Name compression example:

- Query = demo.cs.swarthmore.edu => 4 demo 2 cs 10 swarthmore 3 edu 0
- Response (say NS record) = edu => 20 (pointer to where edu occured previously)
  - rather than encode edu as 03 edu 0, DNS will return a pointer to where edu previously occurred in the packet.

DNS name compression

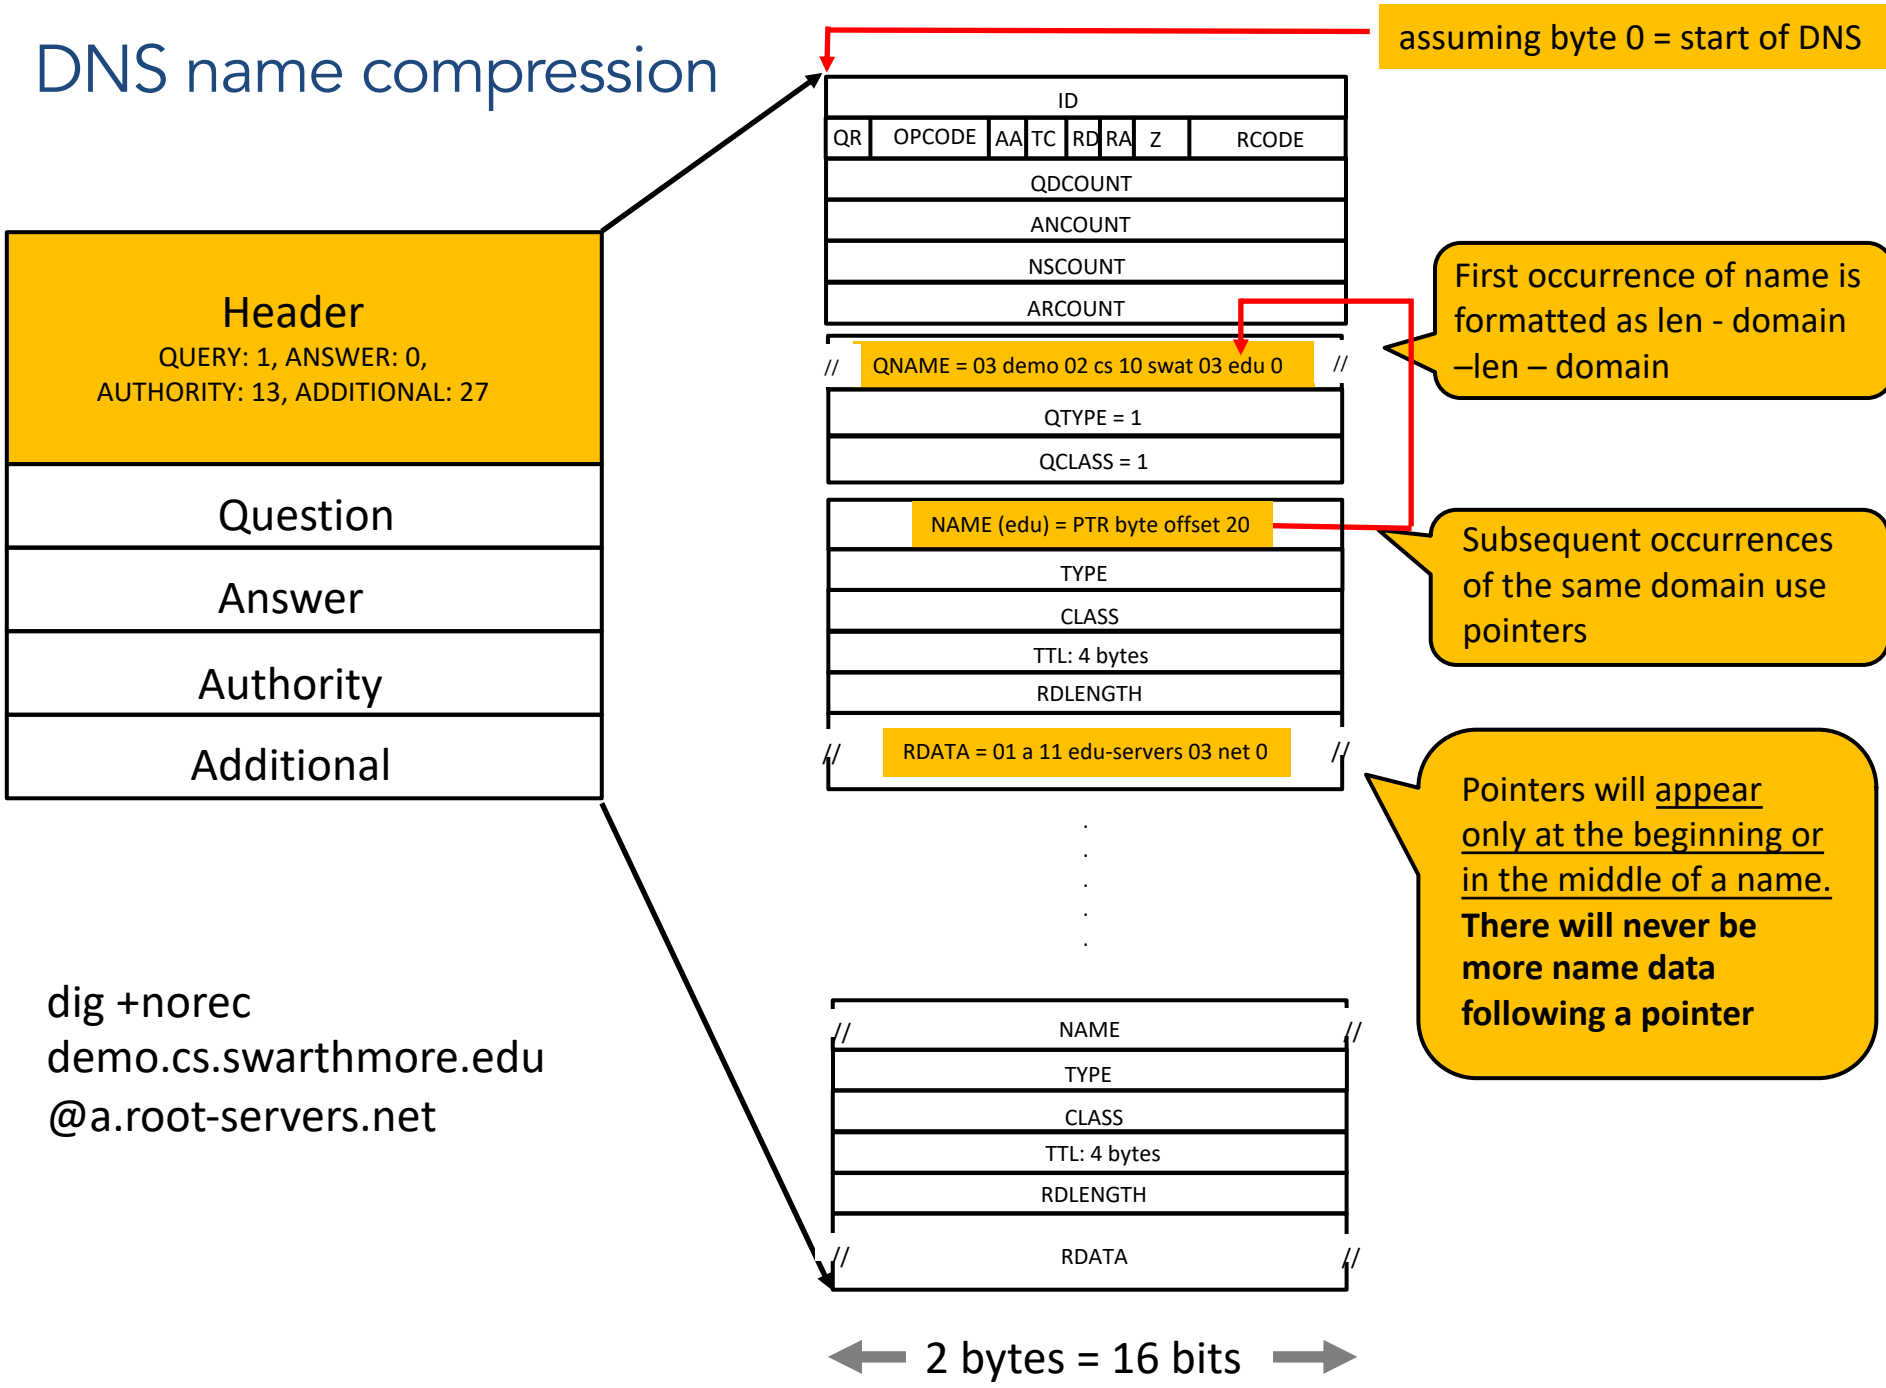

dig +norec
demo.cs.swarthmore.edu
@a.root-servers.net

DNS Name Compression

Is it a name or a pointer?

How can you tell what's contained in a variable length name field?

If it is a variable length name:

- a name either follows a <len><sub-domain> format,
- unpacking the first byte would give you the length of the name field that follows.
- you will know that you've reached the end of the name, when the <len> field is zero (length of the nameless root).

DNS Name Compression

Is it a name or a pointer?

How can you tell what's contained in a variable length name field?

If there is a pointer:

- From the RFC a pointer, is a 2-bytes field, 16 bits in length, and has the format shown below:

```
ascii art format:
index:   0  1  2  3  4  5  6  7  8  9  10 11 12 13 14 15
        +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
        | 1| 1|                  |                       |
        +--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
```

- i.e, first two bits are 1 1
- the next 14 bits specify the offset (from the beginning of the response).

DNS Name Compression

Is it a name or a pointer?

If there is a pointer:

- From the RFC a pointer, is a 2-bytes field, 16 bits in length, and has the format shown below:

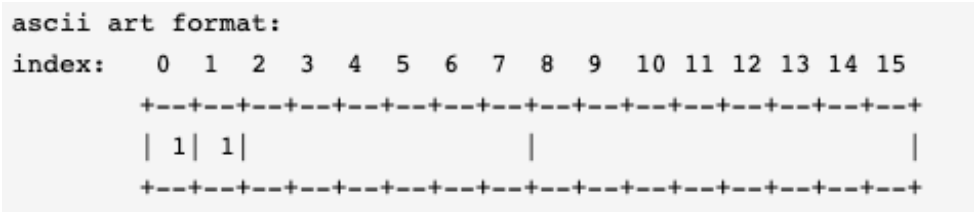

- i.e, first two bits are 1 1
- the next 14 bits specify the offset (from the beginning of the response).

- to figure out if the first byte is a pointer:
  - use *bitwise operations (& |) to flip individual bits such that you return 11000000 only if the first two bits of the first byte are are 11.*
- if it is a pointer, extract the entire 2 byte pointer
  - to get the offset, you'll again need bitwise operations *to only return the 14 bits and exclude the first two bits.*

DNS fields you should be able to parse

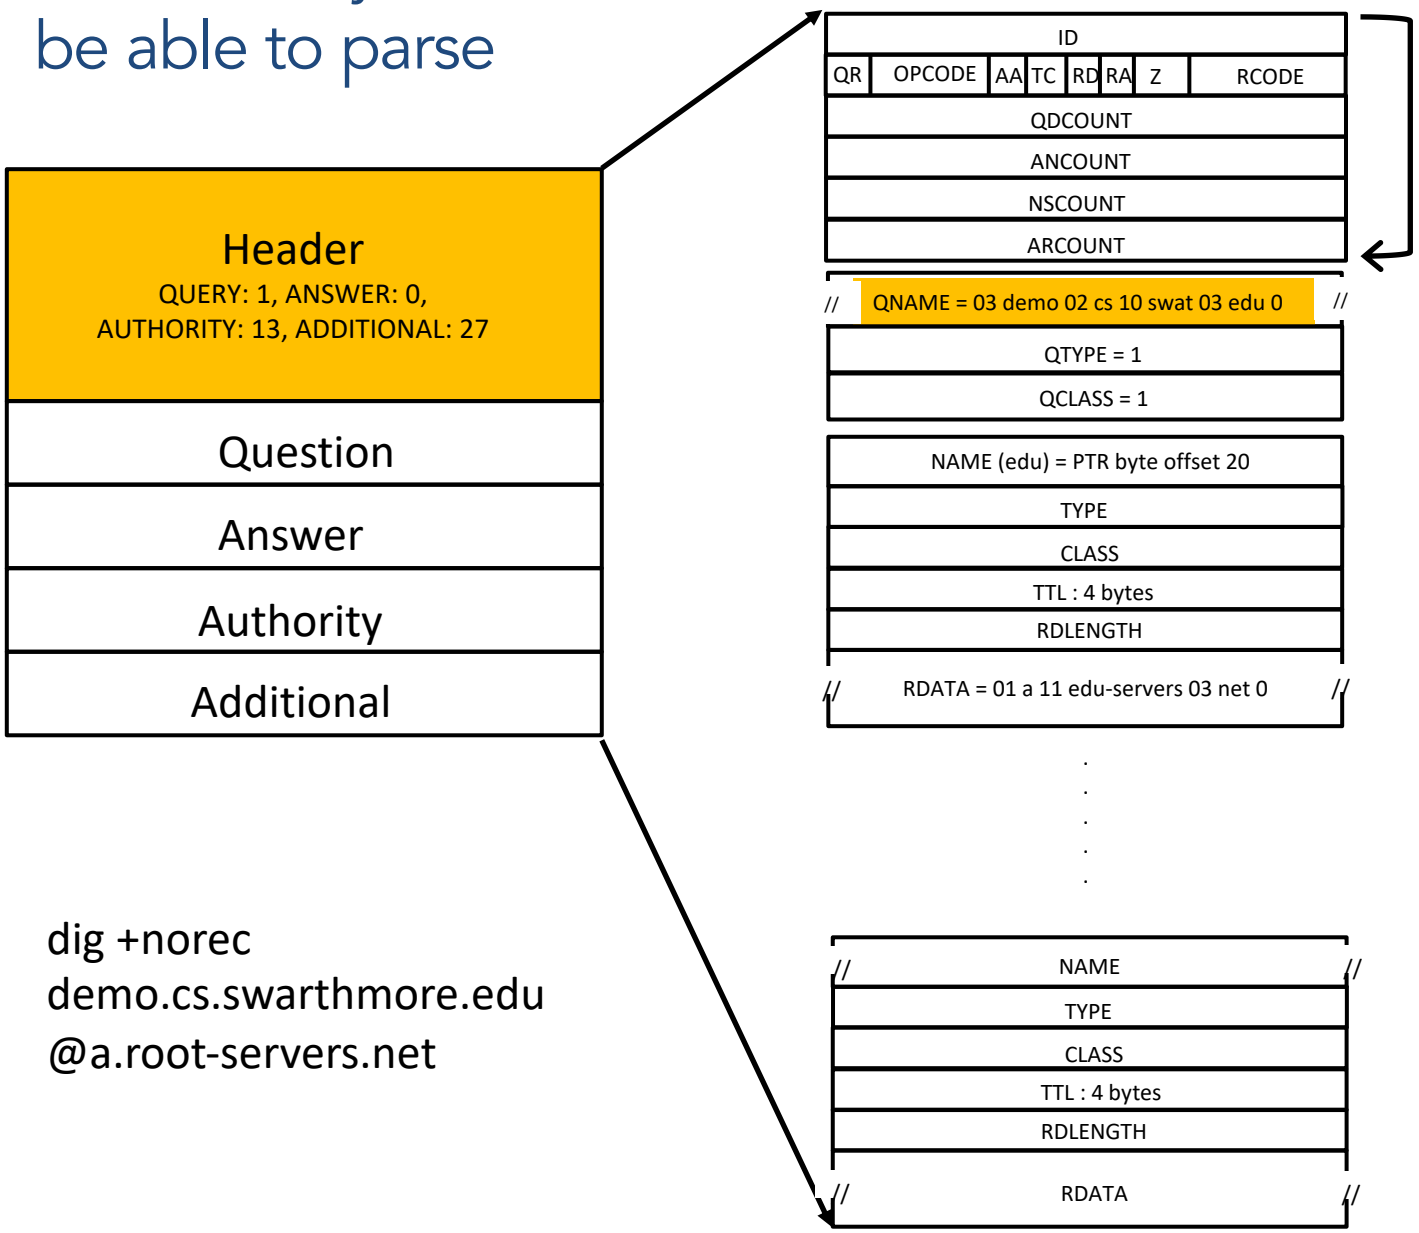

Fixed length fields

Variable length fields

Keep an index of where you are at in the DNS response

dig +norec
demo.cs.swarthmore.edu
@a.root-servers.net

DNS fields you care about

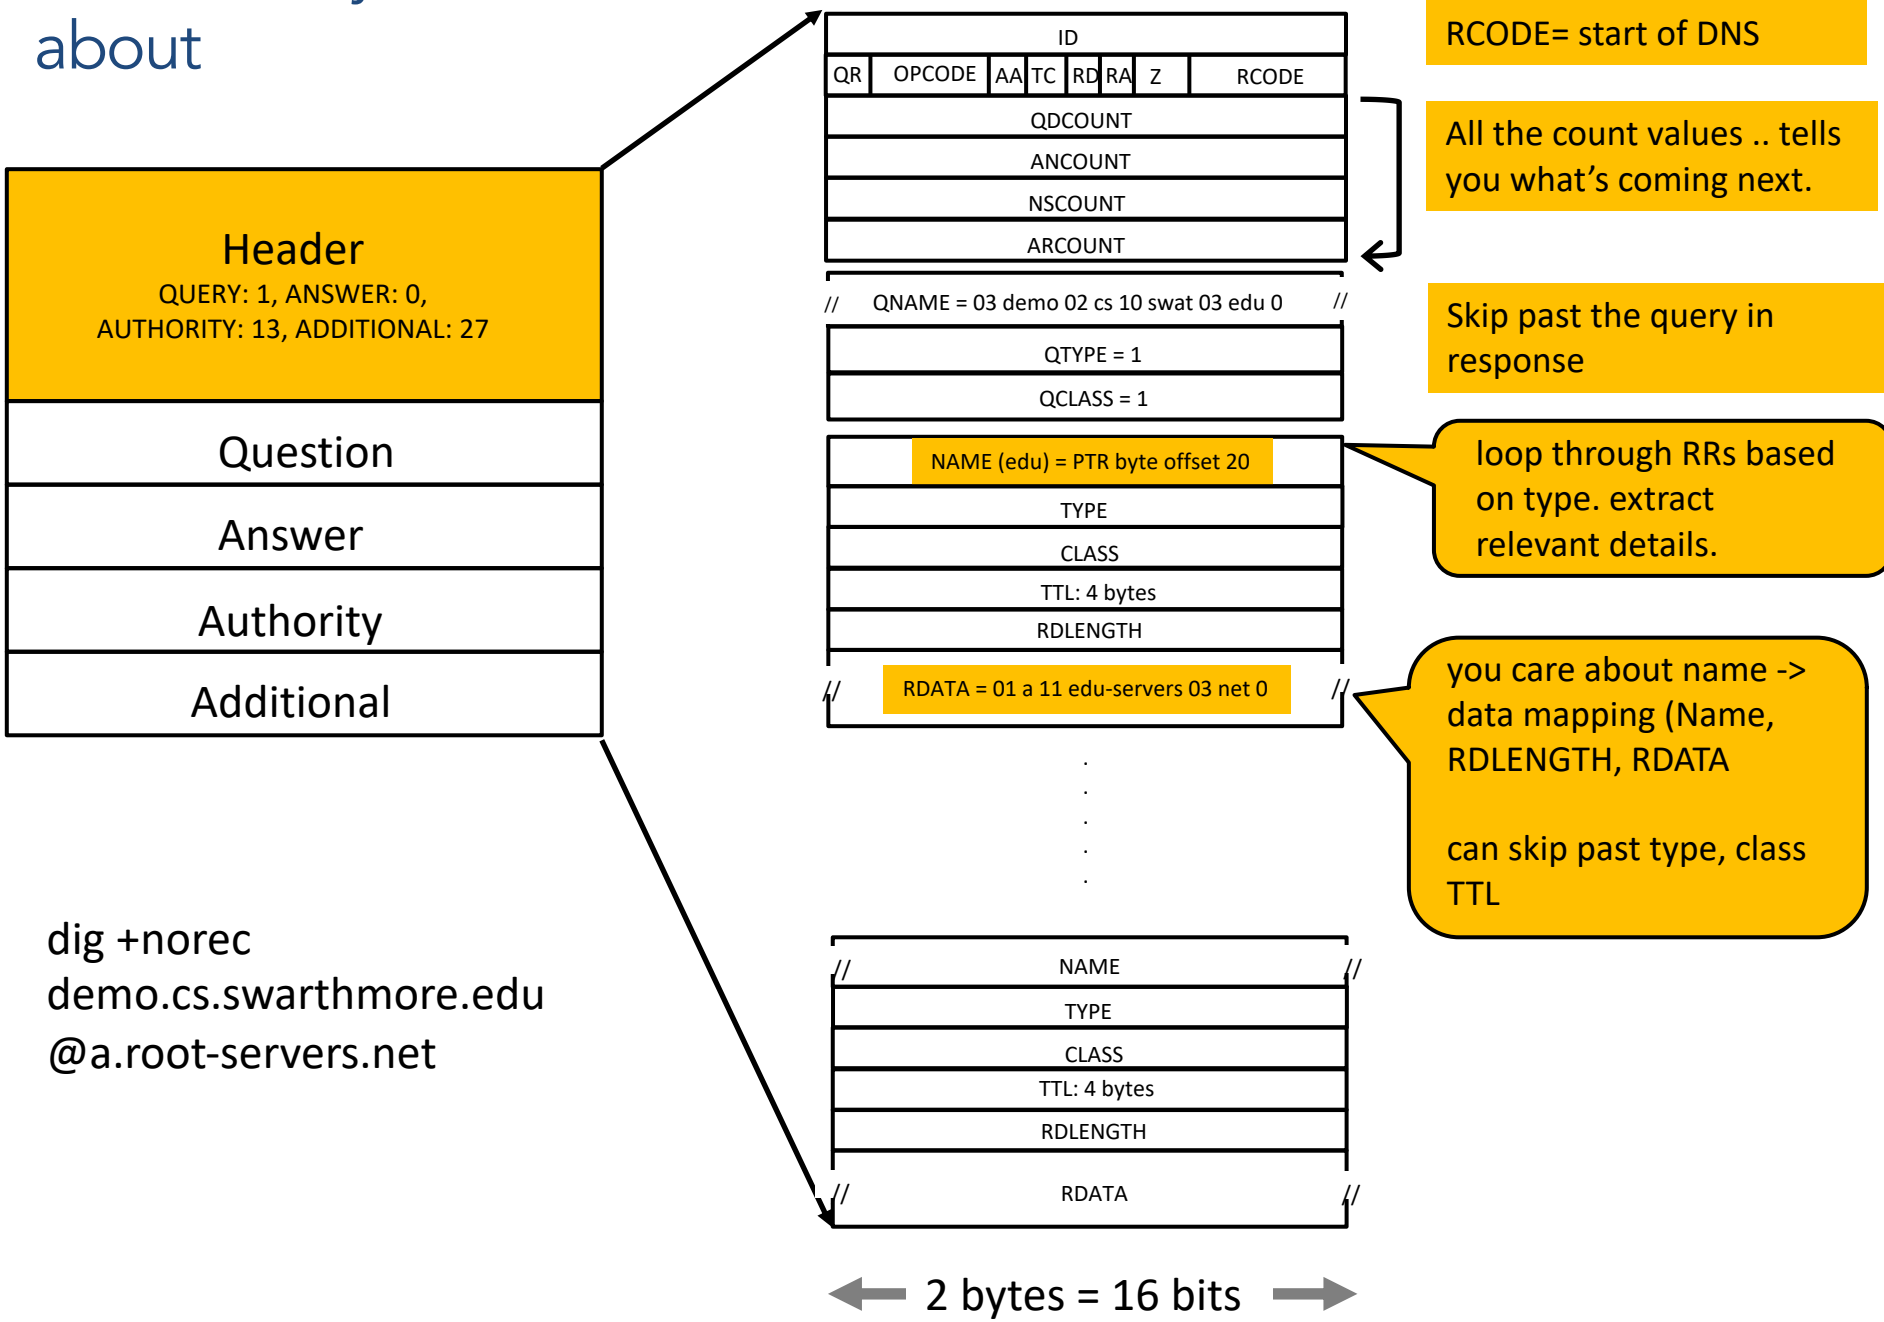

dig +norec
demo.cs.swarthmore.edu
@a.root-servers.net