The Domain Name System

DNS

- This is the means by which we can convert names like news.bbc.co.uk into IP addresses like 212.59.226.30
- Purely for the benefit of human users: we can remember numbers (e.g., telephone numbers), but find it much easier to remember names

The Domain Name System

DNS

- In the early Internet every host had a copy of every machine's name an corresponding IP address: this is no longer possible
- DNS developed as a hierarchical system to *resolve* names and is *distributed*
- No single machine anywhere has all the names of all the machines on the Internet
- The mapping is spread over a large number of hosts

The Domain Name System

The Hierarchy

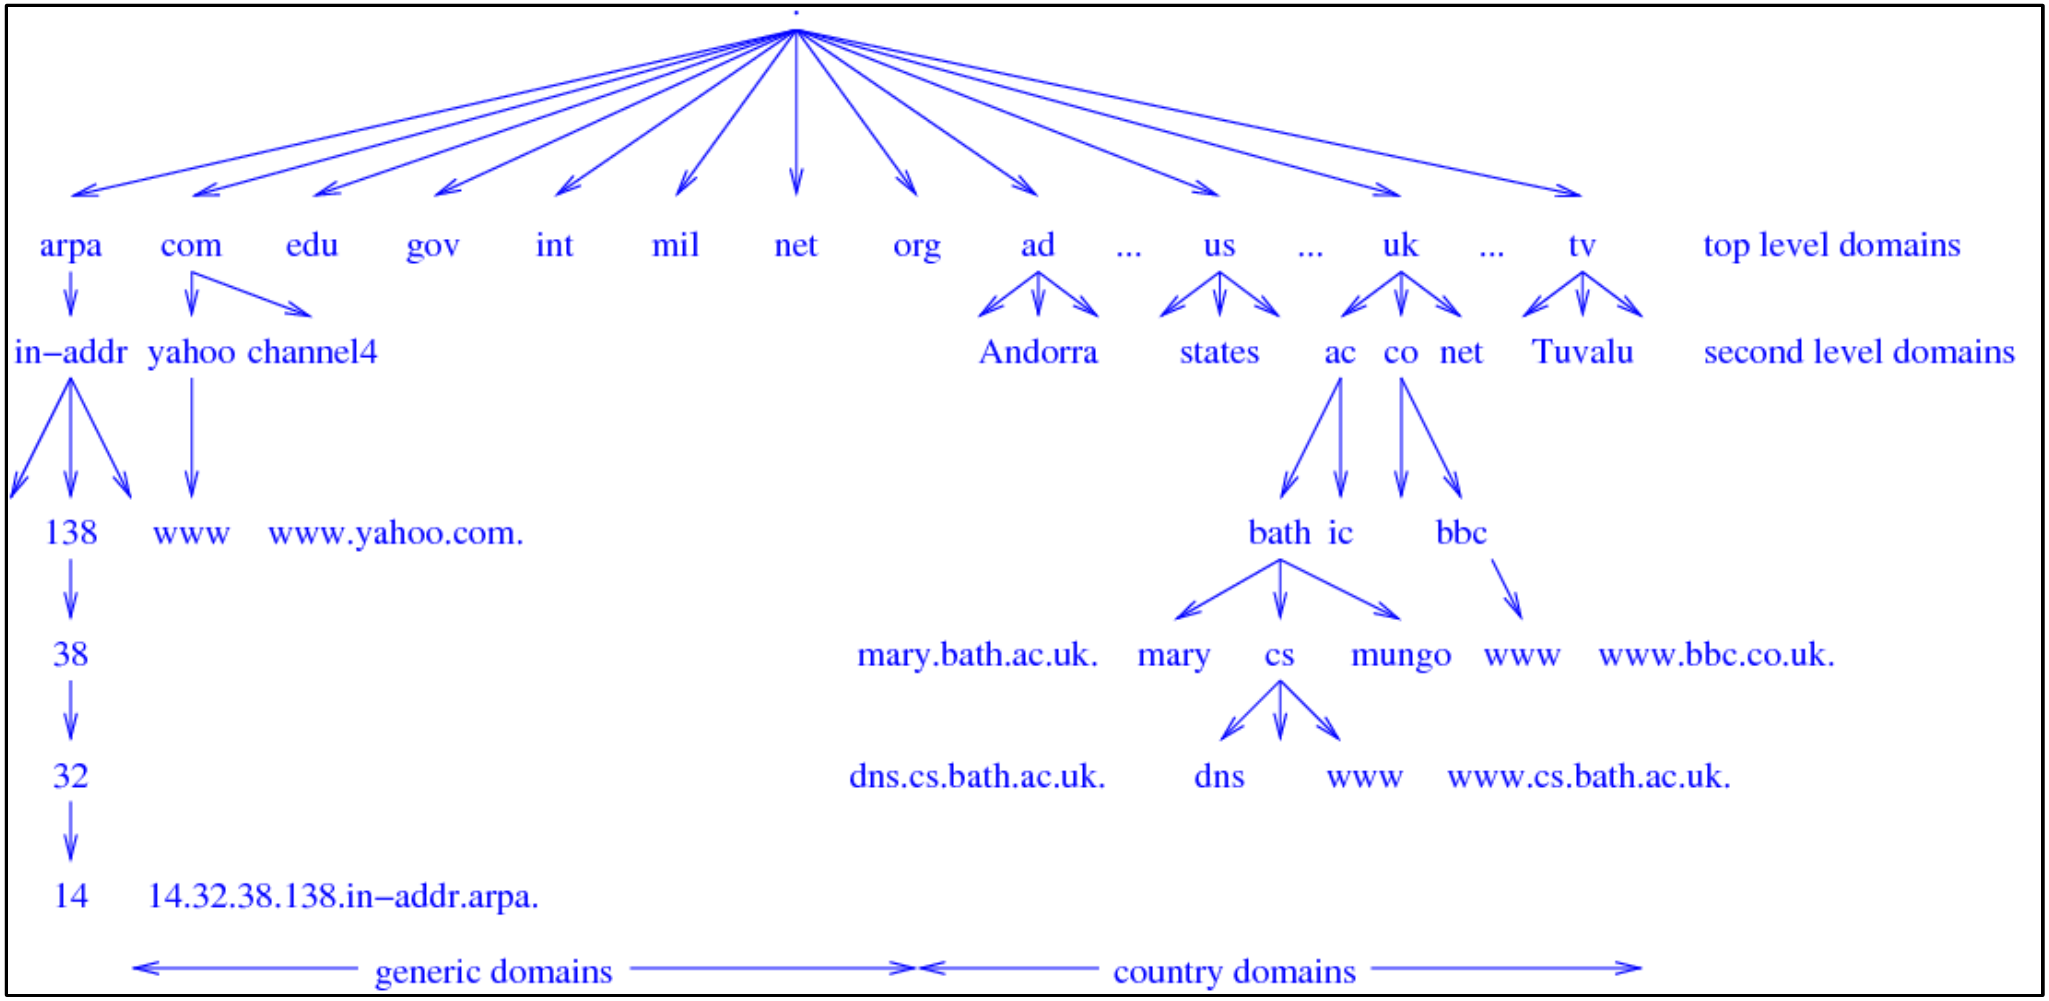

The Domain Name System

The Hierarchy

- A tree with the root at the top, named "." (dot)
- Other nodes in the tree have *labels* of up to 63 characters
- A *fully qualified domain name* (FQDN) is a sequence of labels terminated by a dot, e.g., www.bbc.co.uk.

The Domain Name System

FQDN

- A name without a terminating dot is incomplete, and must be completed before we look up the IP address
- Usually the name is completed by the system appending strings as found in files like /etc/resolv.conf
- Example: mary would be completed to mary.bath.ac.uk. (note the final dot)

The Domain Name System

Hierarchy

- bath.ac.uk is the *network domain*
- For a long time there were just seven *generic domains*, each three characters long
- These included com, org and edu
- Mostly for machines in the USA, but some also worldwide
- Recently more have been added, including biz and info

The Domain Name System

Hierarchy

- Two character names are *country domains* and refer to the relevant country
- From ISO3166, the official list of country abbreviations
- Except the UK, who use uk instead of the ISO gb

The Domain Name System

Hierarchy

- Each level in the tree represents a different management and responsibility for the names
- The *top level domains* (TLDs) are managed by IANA, delegated to ICANN
- Other nodes are managed by other entities

The Domain Name System

Hierarchy

- uk is managed by the Nominet company
- ac.uk is managed by the United Kingdom Education and Research Networking Association (UKERNA)
- co.uk is also managed by Nominet
- bbc.co.uk is managed by the BBC

The Domain Name System

Hierarchy

- Similar names can resolve to completely unrelated IP addresses
  - bill.acme.com might be in Rangoon
  - ben.acme.com might be in Tunbridge Wells
- tv belongs to the island of Tuvalu, but they have sold the rights to the name to places all over the world, e.g., bbc.tv

The Domain Name System

Hierarchy

- Hierarchy allows structure

ac.uk	Academic establishments	net.uk	Internet service providers
bl.uk	British Library	nhs.uk	National Health Service
co.uk	Commercial enterprises	org.uk	Not-for-profit organisations
gov.uk	Government bodies	plc.uk	Public limited companies
ltd.uk	Limited companies	police.uk	Police
me.uk	Personal domains	sch.uk	Schools
mod.uk	Ministry of Defence		

The Domain Name System

Hierarchy

- A *zone* is a subtree that is administered separately, e.g., bath.ac.uk
- Zones can have subzones, e.g., cs.bath.ac.uk
- The authority for a zone must set up *name servers*. A name server is (a program on) a machine that has the database of labels for that zone

The Domain Name System

Servers

- Names are added or deleted at this level by changing this database
- For resilience, there must be a *primary name server*, and one or more *secondary name servers*
- The secondaries get their copy of the data from the primary by periodic *zone transfers*, which is just the copying of the database

The Domain Name System

Servers

- It is good practice to have a secondary off-site and even off-continent to increase resilience even further
- If a machine requests a name lookup from a server that is in its own zone, the server can provide an *authoritative* reply
- Otherwise lookup is harder and this is the main point of the DNS

The Domain Name System

Recursive Lookup

- If a request is for a name outside the authority of a server, the server must ask around for the name
- It first contacts a *root name server*: this is one of currently about 80 servers spread about the world that are responsible for the TLDs (the root zone)

The Domain Name System

Recursive Lookup

- Using load sharing these machines share just 13 IP addresses and are named a.root-servers.net to m.root-servers.net
- The IP addreses of these machines are stored locally on our name server: otherwise we could never get started!
- All the root servers contain the same information: this spreads load and provides resilience against attack

The Domain Name System

```
;       This file holds the information on root name servers needed to
;       initialize cache of Internet domain name servers
;       (e.g. reference this file in the "cache  .  <file>"
;       configuration file of BIND domain name servers).
;
     ...
; formerly NS.INTERNIC.NET
;
.                       3600000  IN  NS    A.ROOT-SERVERS.NET.
A.ROOT-SERVERS.NET.      3600000      A     198.41.0.4
;
; formerly NS1.ISI.EDU
;
.                       3600000      NS    B.ROOT-SERVERS.NET.
B.ROOT-SERVERS.NET.      3600000      A     192.228.79.201
;
; formerly C.PSI.NET
;
...
; operated by WIDE
;
.                       3600000      NS    M.ROOT-SERVERS.NET.
M.ROOT-SERVERS.NET.      3600000      A     202.12.27.33
; End of File
```

The Domain Name System

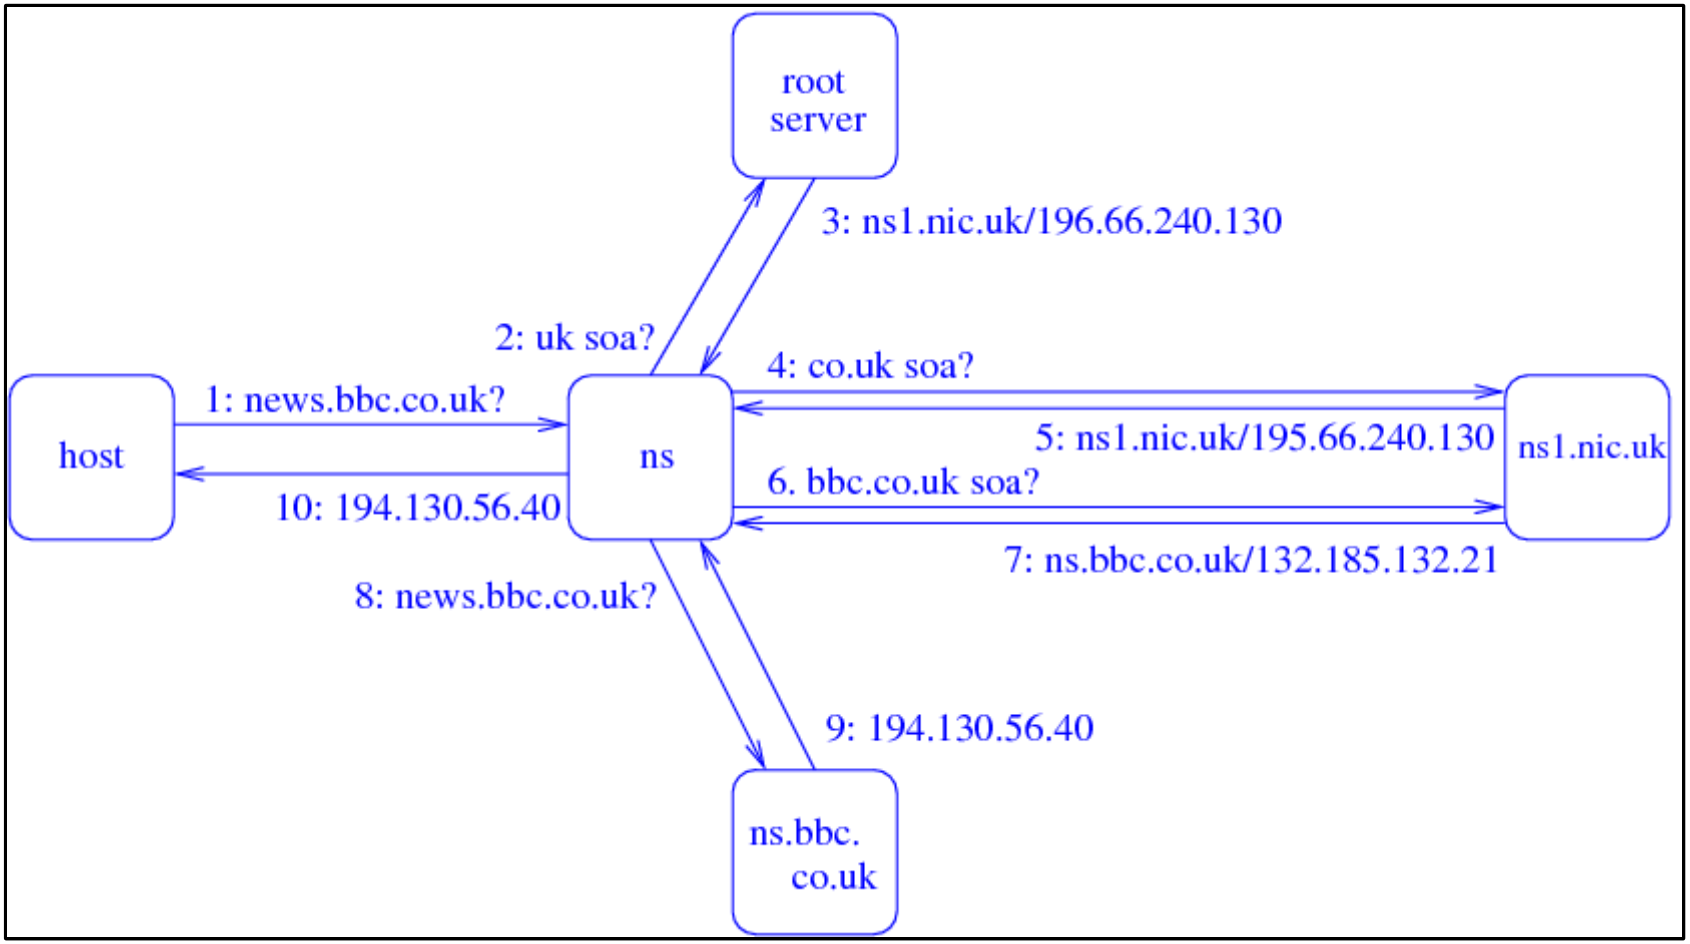

1. We want to find the IP address of news.bbc.co.uk, so we ask our name server ns

The Domain Name System

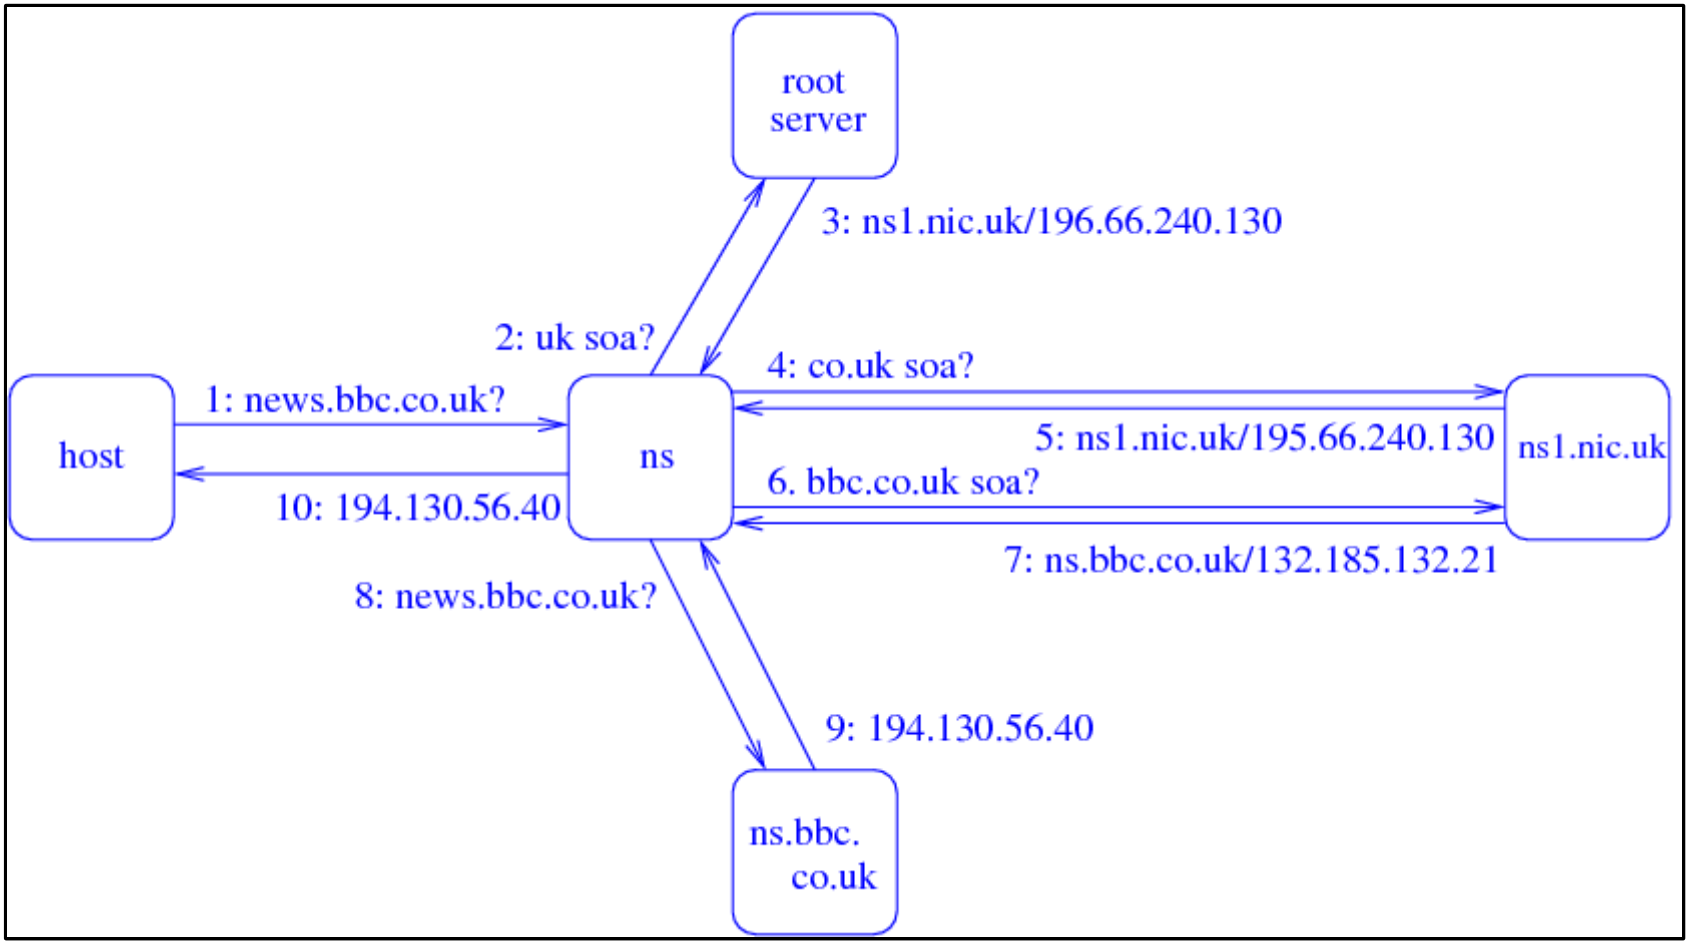

- Our name server ns does not have responsibility for the bbc.co.uk domain

The Domain Name System

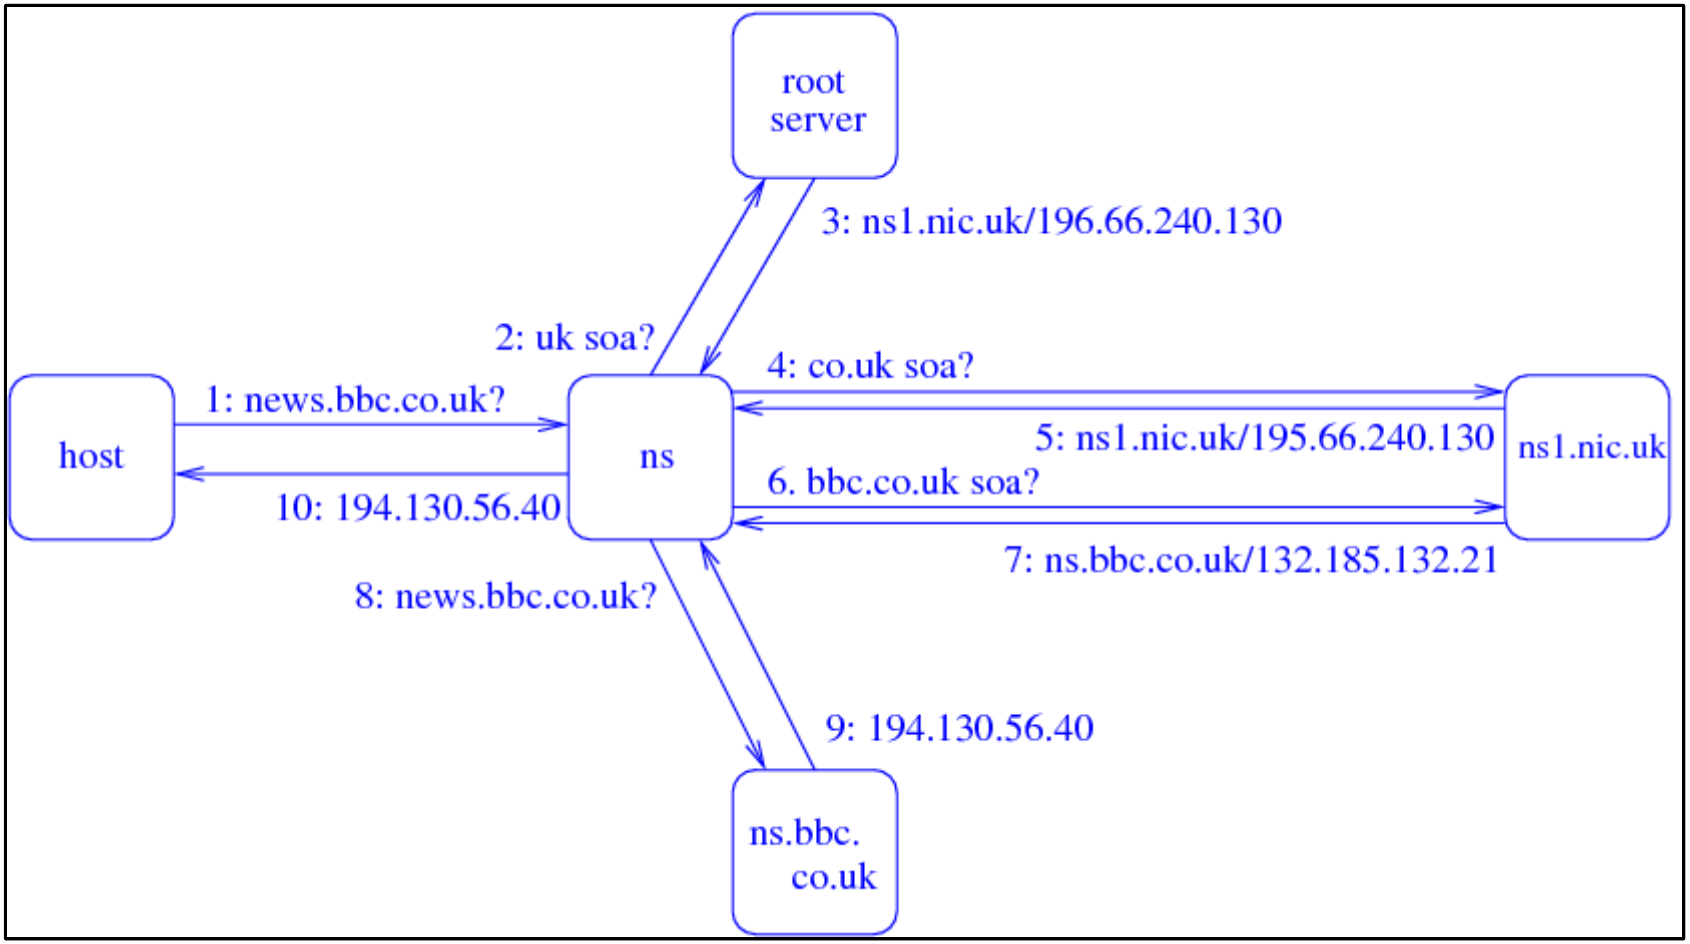

2. So it asks a random root server: “who has responsibility for the uk domain?” This is a *start of authority* (SOA) request

The Domain Name System

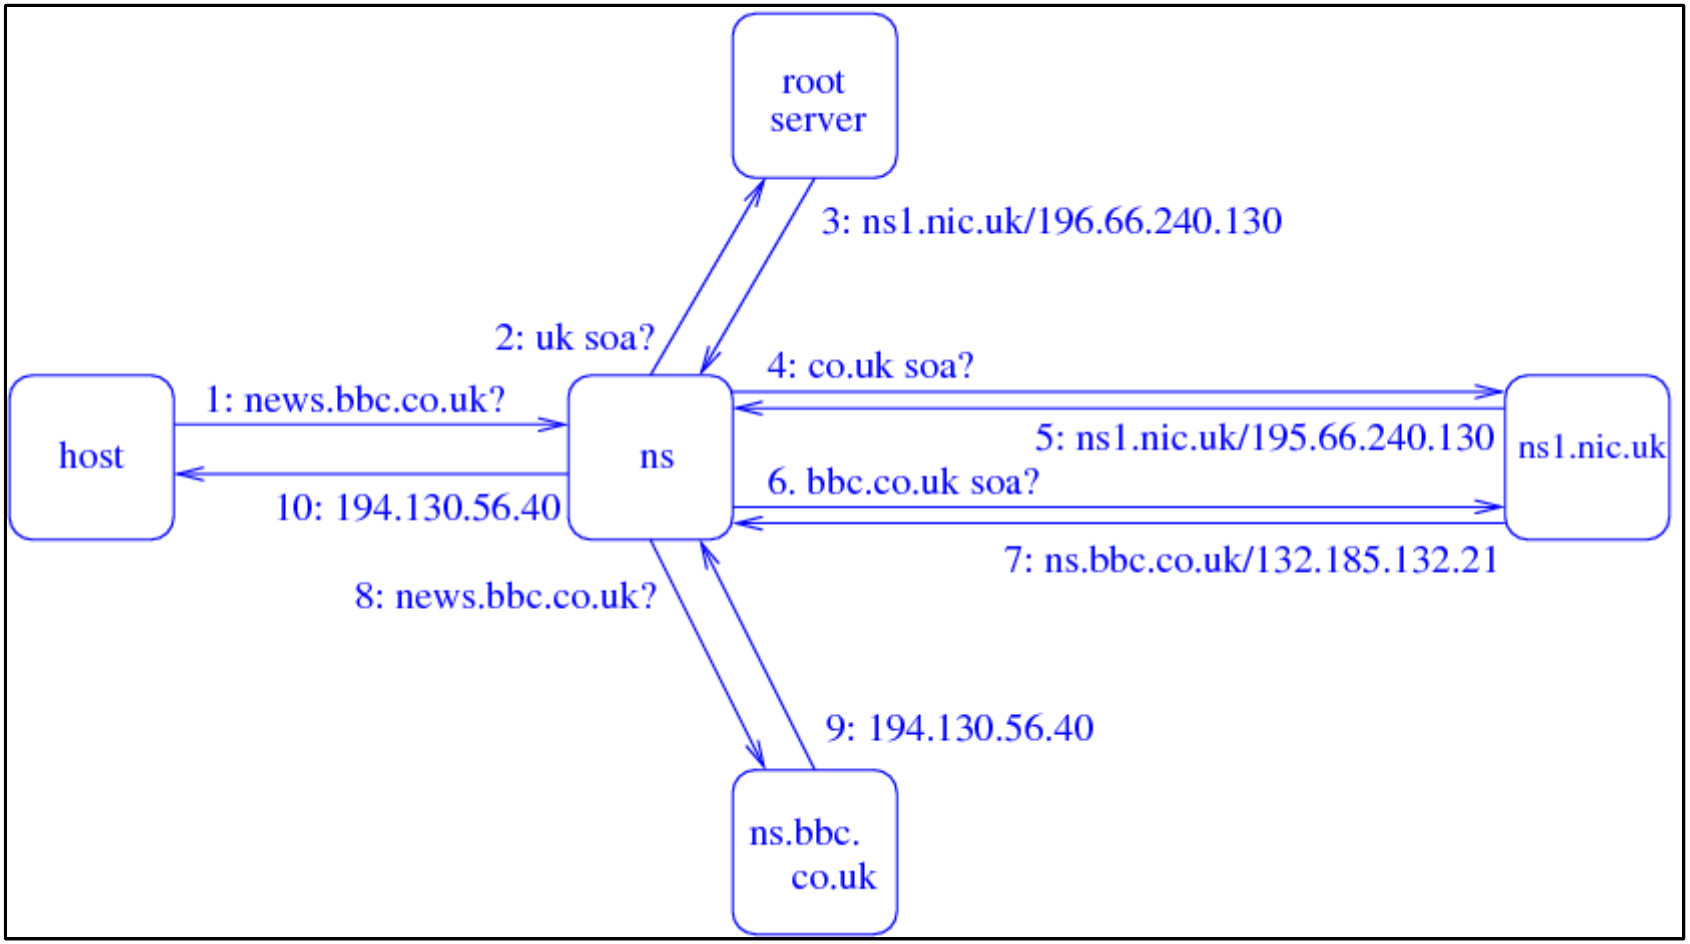

3. The server replies ns1.nic.uk and provides the IP address of that machine

The Domain Name System

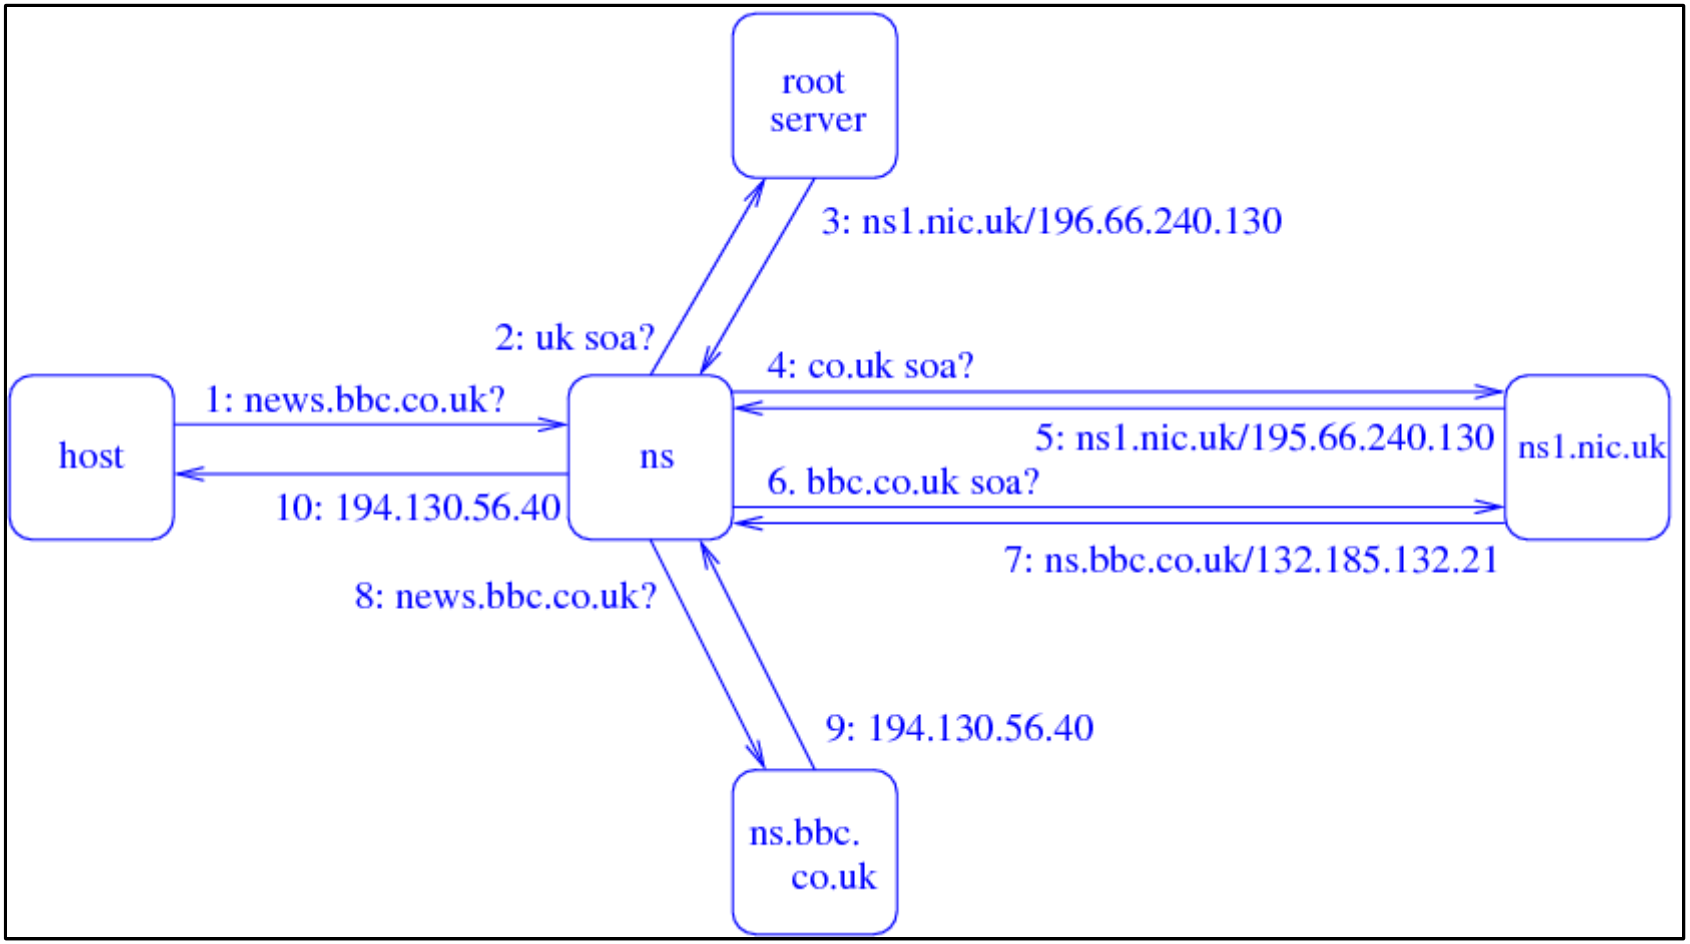

4. ns now asks ns1.nic.uk “who has responsibility for co.uk?”

The Domain Name System

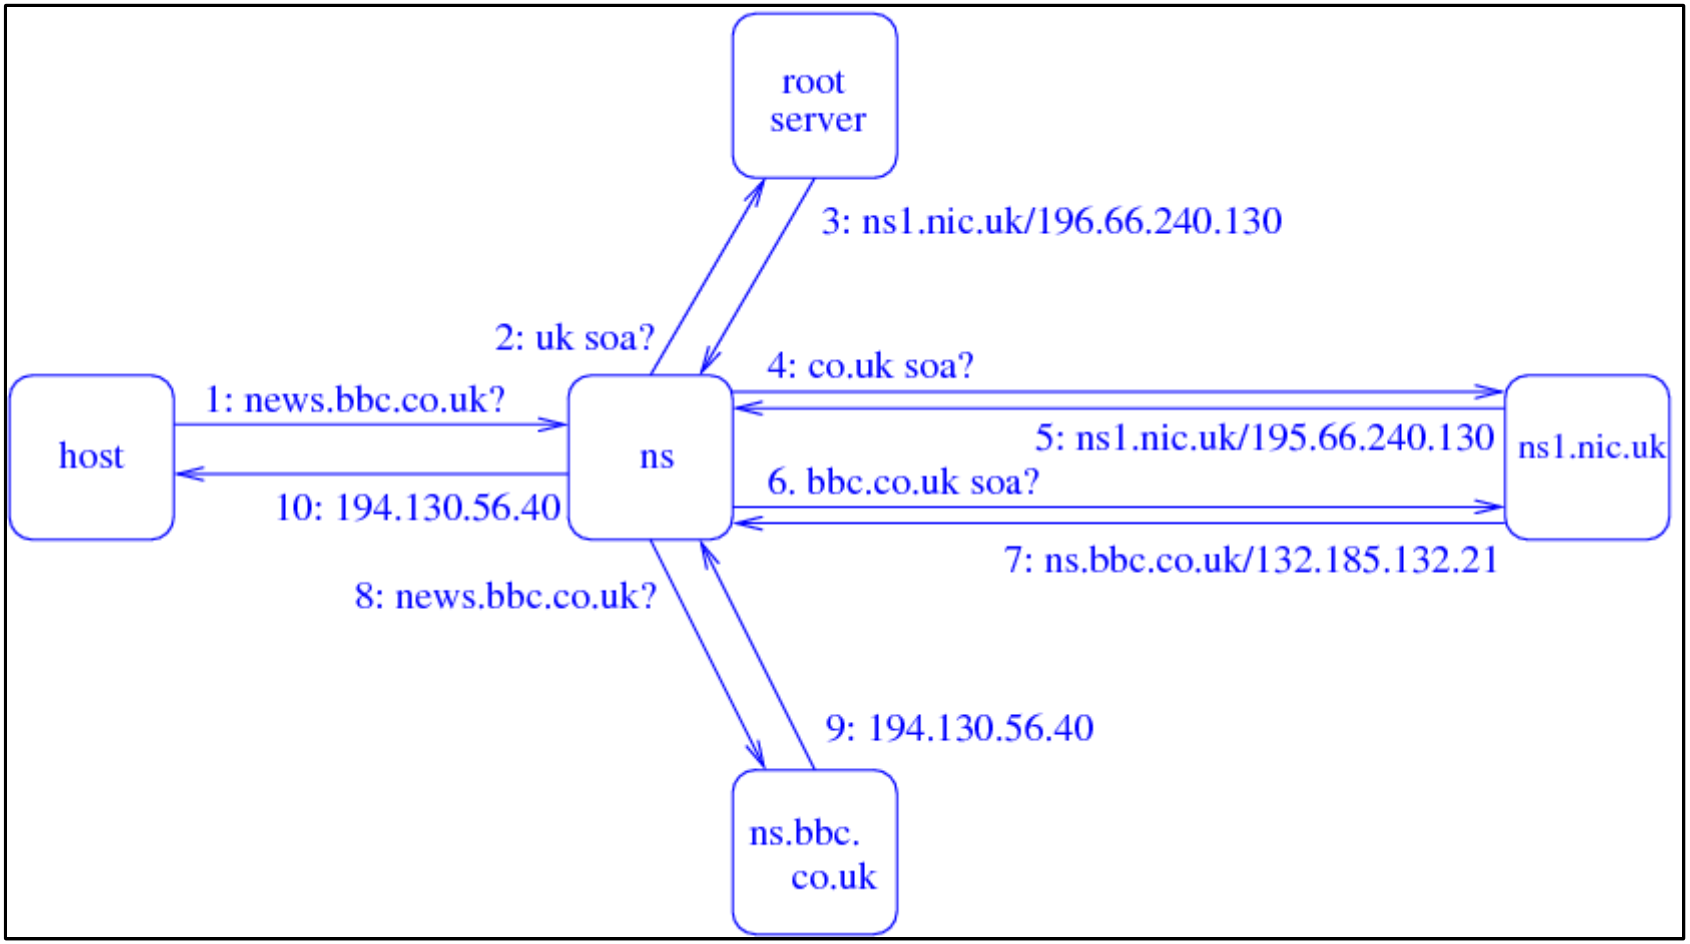

5. It replies ns1.nic.uk. It just so happens it is also the SOA for co.uk

The Domain Name System

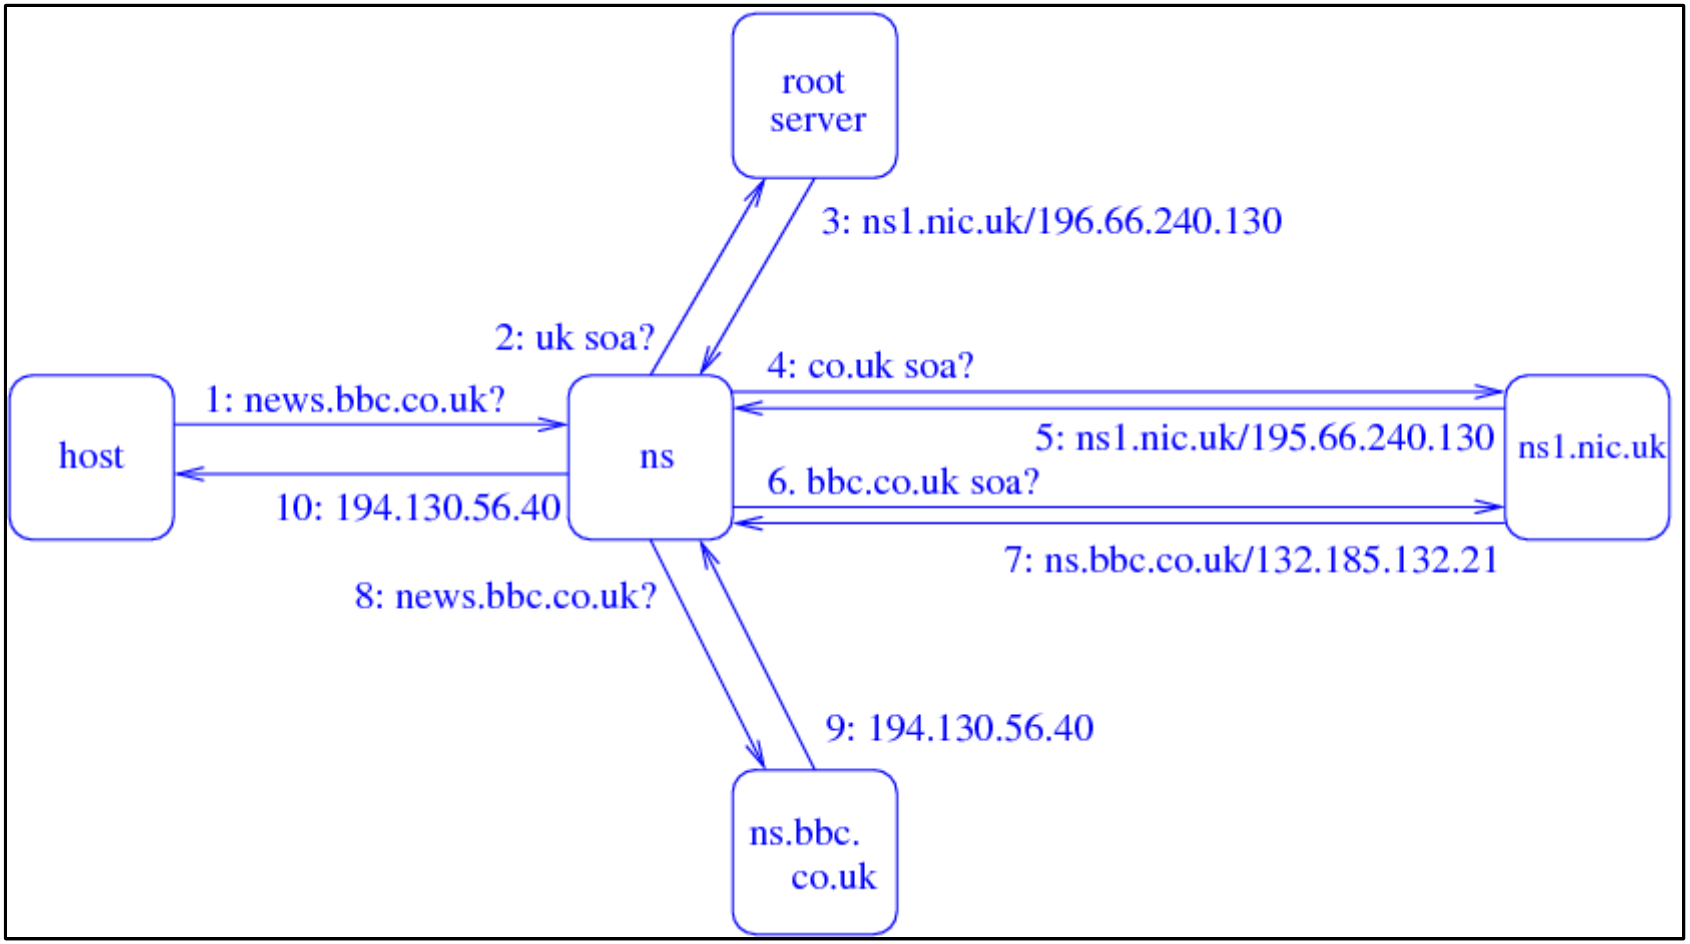

6. ns asks ns1.nic.uk “who is responsible for bbc.co.uk?”

The Domain Name System

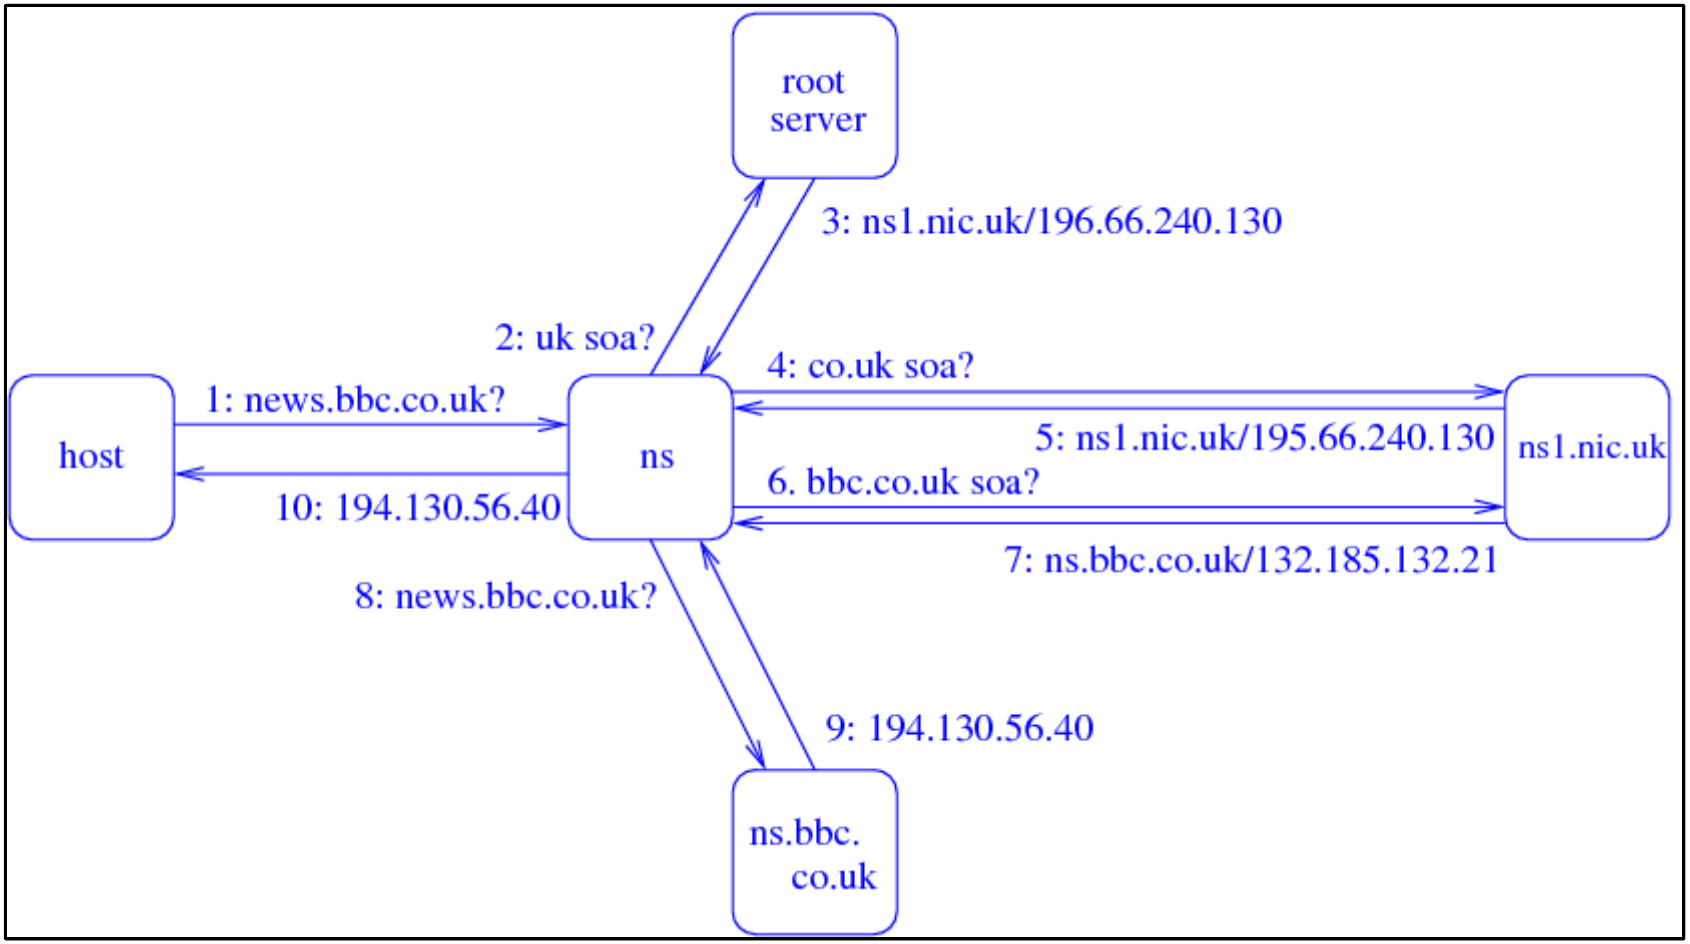

7.It replies ns.bbc.co.uk

The Domain Name System

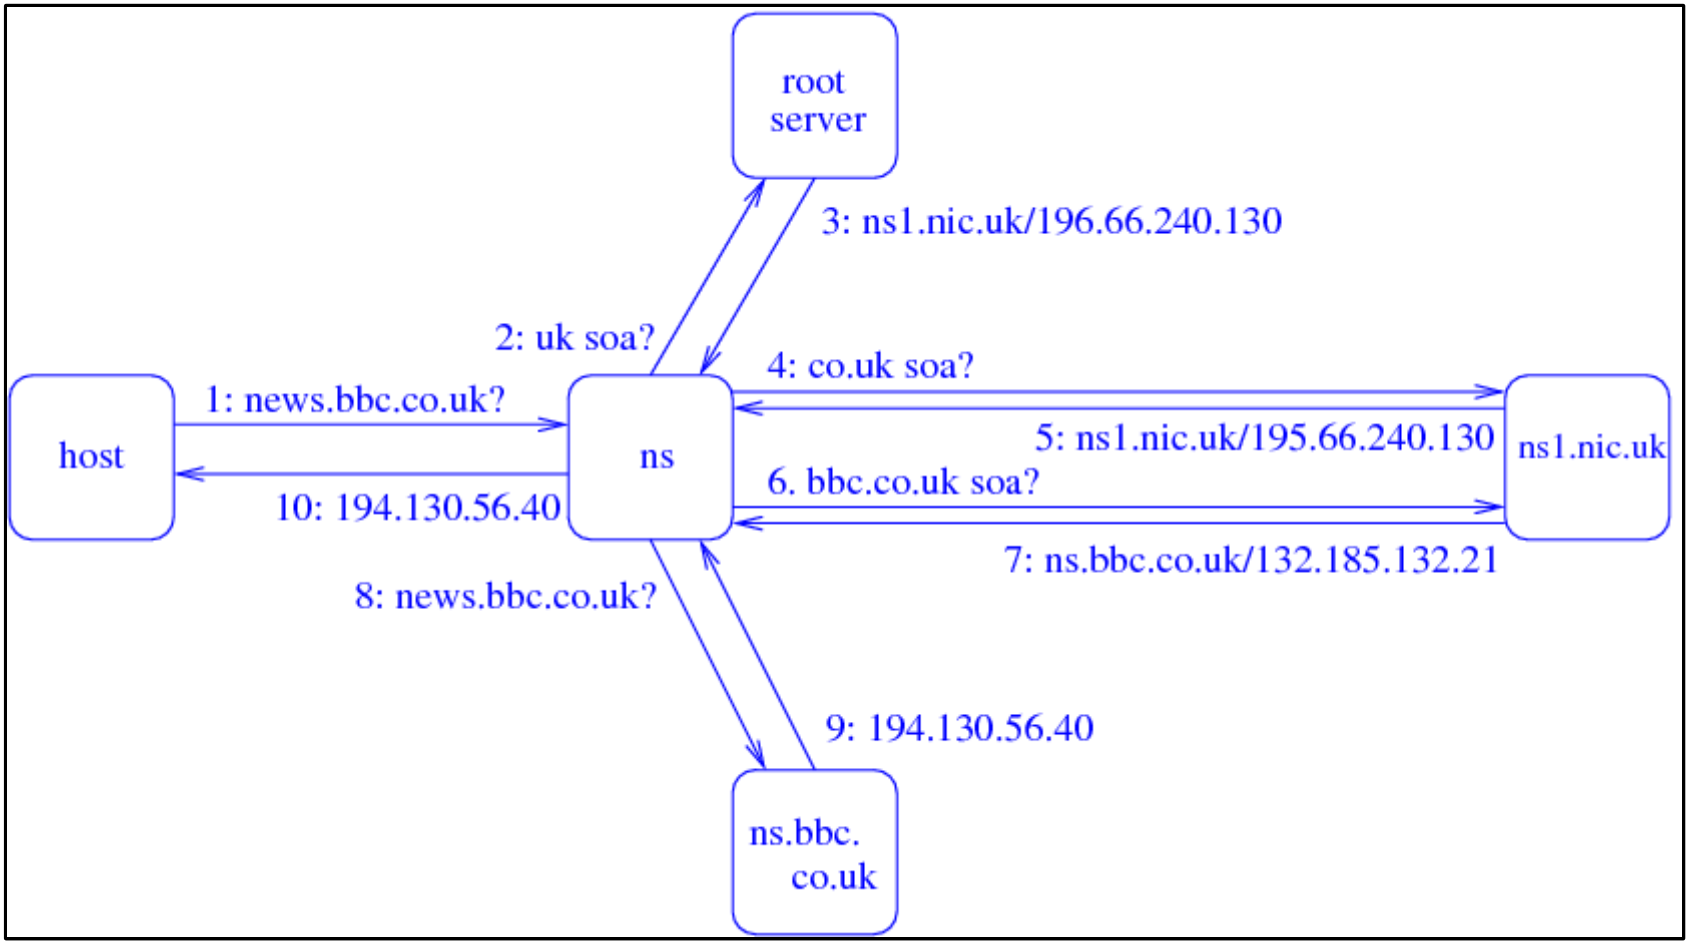

8. We have now found the authority for the domain we are looking for. ns asks ns.bbc.co.uk “what is the IP address of news.bbc.co.uk?”

The Domain Name System

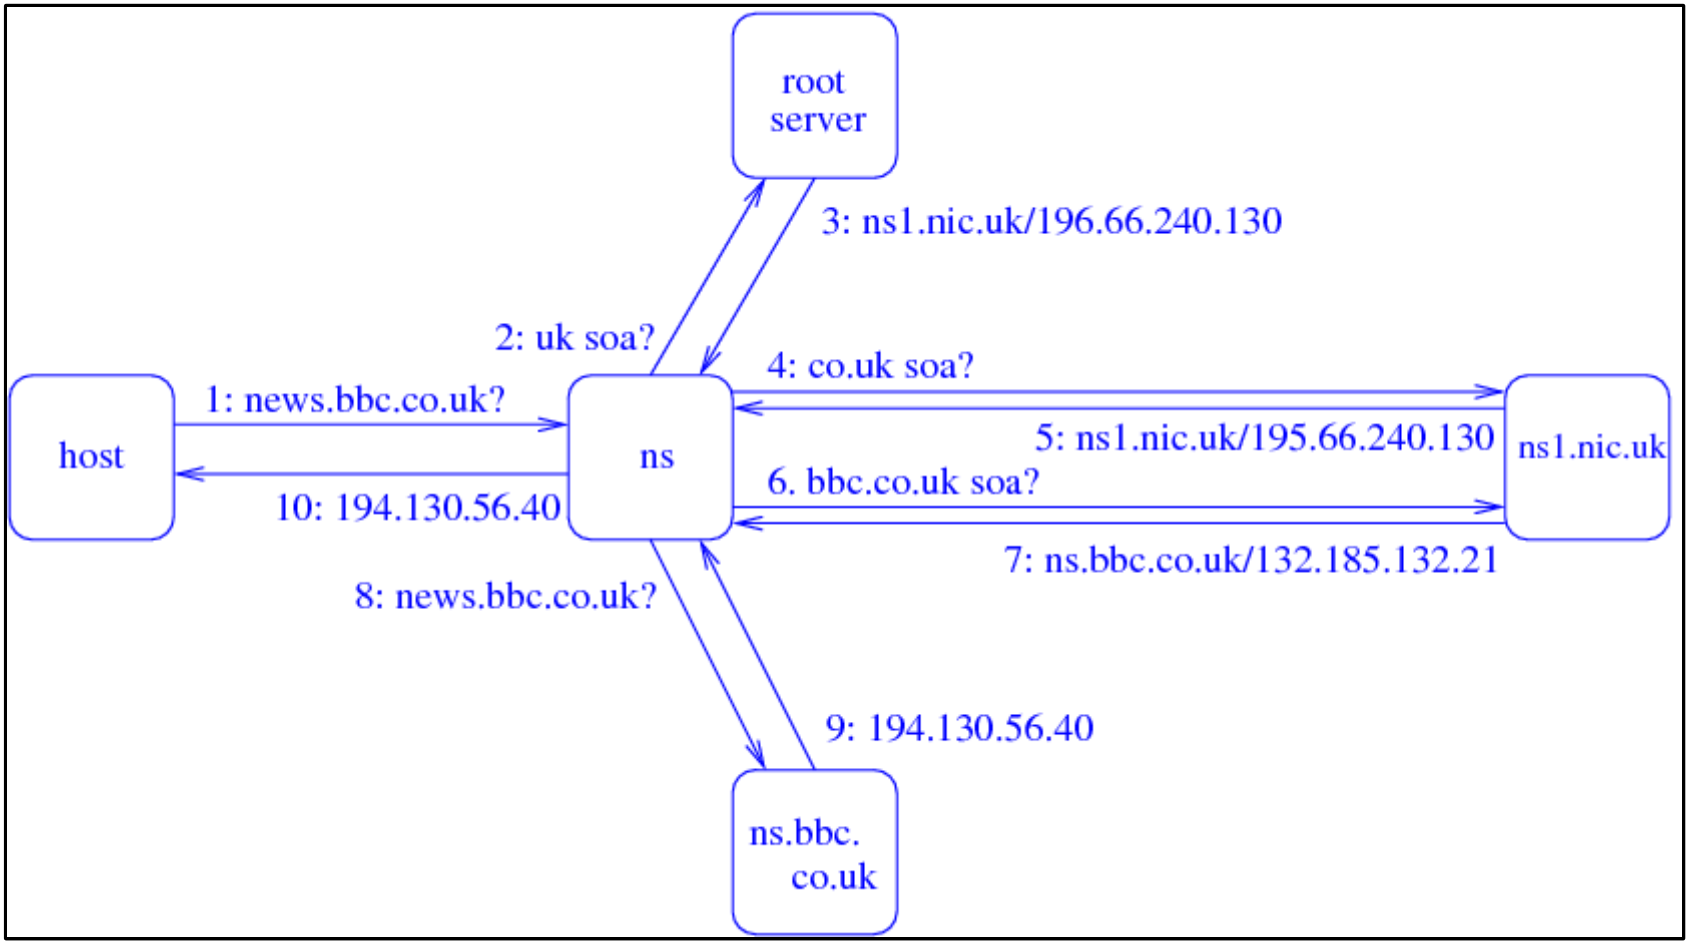

9.It replies with the IP address

The Domain Name System

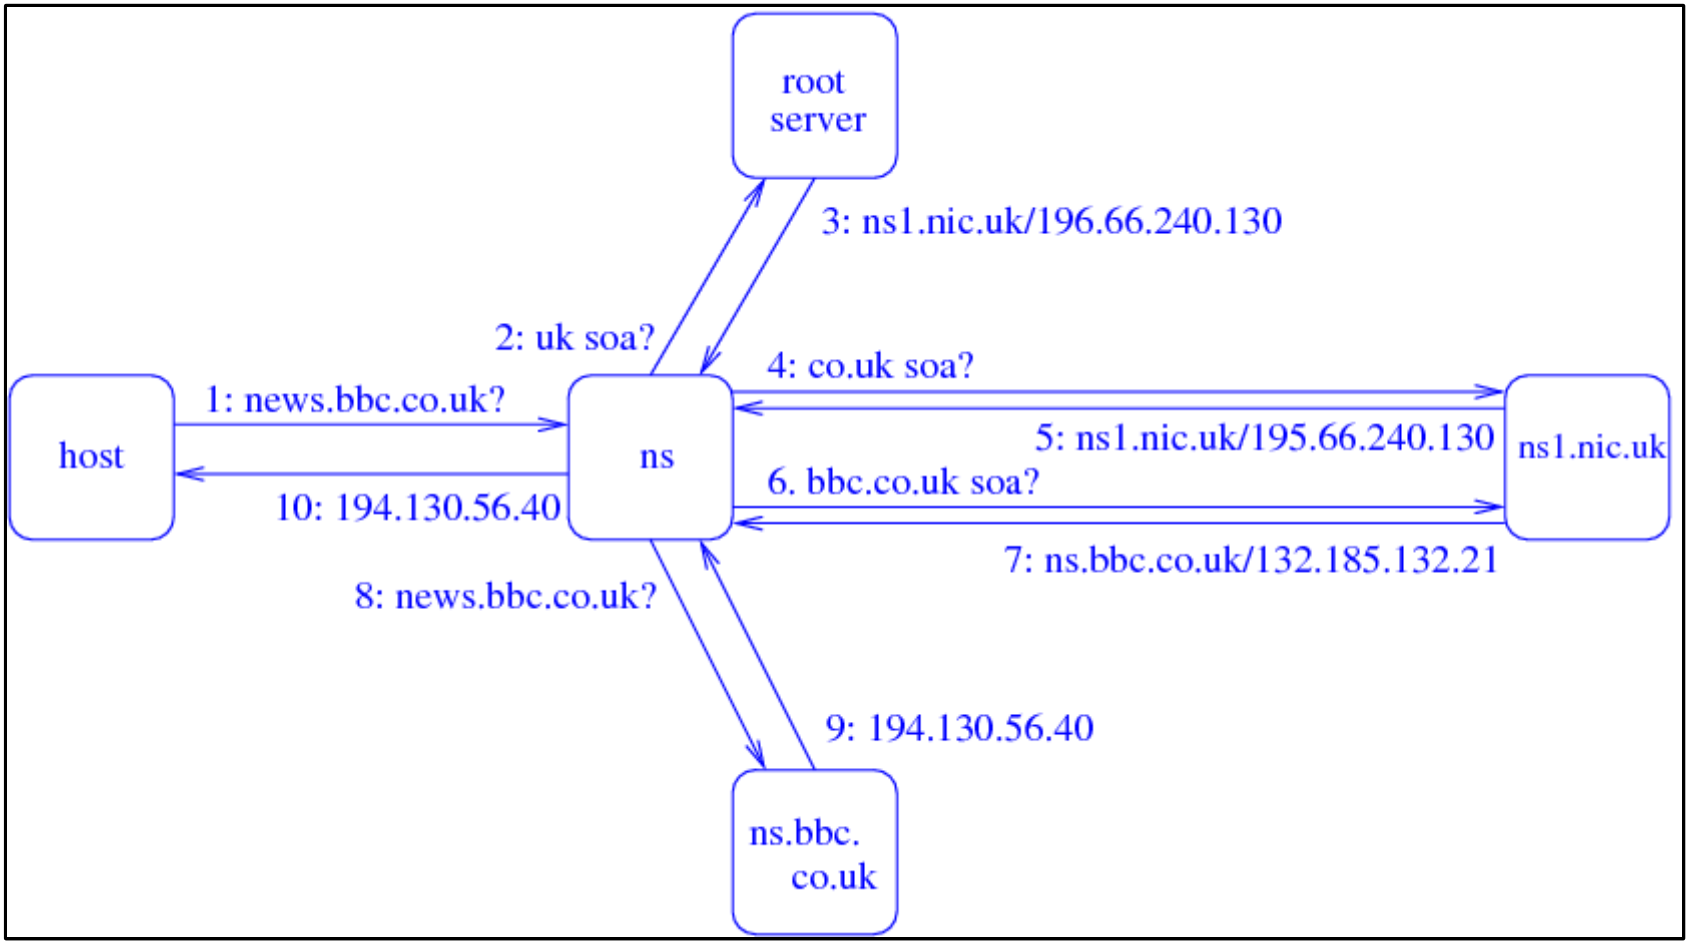

10. Finally, ns relays the answer back to the host that made the original request

The Domain Name System

Recursive Lookup

- ns has done a *recursive lookup* on behalf of the original host
- All the responses are cached by ns so that is does not have to go through the full process every time a lookup is requested
- Each reply has a *time to live* attached that indicates how long the server should keep the information before asking again

The Domain Name System

Recursive Lookup

- The next request for a bbc.co.uk name can go directly to ns.bbc.co.uk without bothering ns1.nic.uk or a root server
- The host could do the lookup process itself (a non-recursive lookup), but an advantage of caching in the name server means any host in the organisation can benefit from other hosts' lookups, thus speeding up the response

The Domain Name System

DNS

- The distributed nature of the DNS is a great strength against attack
- In October 2002 there was a malicious attempt to flood the root servers with bogus requests from subverted hosts and thus bring the Internet to a halt
- There was 30-40 times the normal DNS traffic

The Domain Name System

DNS

- Barely anyone (other than the root server operators) noticed
- The efficiency of the DNS was reduced to about 94% of normal

The Domain Name System

```
% dig news.bbc.co.uk

...

;; QUESTION SECTION:
;news.bbc.co.uk.                    IN      A

;; ANSWER SECTION:
news.bbc.co.uk.         360     IN      CNAME   newswww.bbc.net.uk.
newswww.bbc.net.uk.     62      IN      A       212.58.226.19

;; AUTHORITY SECTION:
bbc.net.uk.             162395  IN      NS      ns0.thdo.bbc.co.uk.
bbc.net.uk.             162395  IN      NS      ns0.thny.bbc.co.uk.

;; ADDITIONAL SECTION:
ns0.thdo.bbc.co.uk.     211     IN      A       212.58.224.20
ns0.thny.bbc.co.uk.     211     IN      A       212.58.240.20

...
```

The Domain Name System

DNS

- More than one name can map on to the same IP address: newswww.bbc.net.uk is the *canonical name* (CNAME) for this machine, while news.bbc.co.uk is an *alias*
- Aliases are useful to give mnemonic names to machines. We can change the machine a web server (say) runs on without changing the name: just change the alias

The Domain Name System

DNS

- The reply also give information about the name servers (NS) and their addresses (A)
- We can get dig to show the details of the recursive lookup:

dig +trace news.bbc.co.uk

```
...
.                137487  IN      NS      e.root-servers.net.
.                137487  IN      NS      f.root-servers.net.
.                137487  IN      NS      g.root-servers.net.
...
.                137487  IN      NS      b.root-servers.net.
.                137487  IN      NS      c.root-servers.net.
.                137487  IN      NS      d.root-servers.net.

...

uk.                 172800  IN      NS      NSC.NIC.uk.
...
uk.                 172800  IN      NS      NSB.NIC.uk.

...

bbc.co.uk.          172800  IN      NS      ns1.thny.bbc.co.uk.
bbc.co.uk.          172800  IN      NS      ns1.thls.bbc.co.uk.
bbc.co.uk.          172800  IN      NS      ns1.thdo.bbc.co.uk.
bbc.co.uk.          172800  IN      NS      ns1.bbc.co.uk.

...

news.bbc.co.uk.       900     IN      CNAME   newswww.bbc.net.uk.
newswww.bbc.net.uk.   294     IN      A       212.58.226.19
bbc.net.uk.         170057  IN      NS      ns0.thny.bbc.co.uk.
bbc.net.uk.         170057  IN      NS      ns0.thdo.bbc.co.uk.
```

The Domain Name System

DNS

- A name can map to more than one IP address

```
www.yahoo.com.          221    IN     CNAME   www.yahoo.akadns.net.
www.yahoo.akadns.net.   26     IN     A       68.142.226.39
www.yahoo.akadns.net.   26     IN     A       68.142.226.41
www.yahoo.akadns.net.   26     IN     A       68.142.226.43
www.yahoo.akadns.net.   26     IN     A       68.142.226.46
www.yahoo.akadns.net.   26     IN     A       68.142.226.49
www.yahoo.akadns.net.   26     IN     A       68.142.226.54
www.yahoo.akadns.net.   26     IN     A       68.142.226.56
www.yahoo.akadns.net.   26     IN     A       68.142.226.32
```

- If this is done, the IP addresses are used in round-robin fashion. This is another way to balance load

The Domain Name System

DNS

- Usually all the machines sharing a name have identical content so it doesn't matter which one you actually use
- Often these machines are distributed around the world to help spread the network load

The Domain Name System

Reverse DNS

- Another branch of the DNS tree is labelled arpa
- This is used for the reverse problem: given an IP address, find a name
- Just as with names, we delegate from the most important part of the IP address first, which is the first part, rather than the last, as in the name. E.g., 138.38.32.14, delegate from the 138

The Domain Name System

Reverse DNS

- Subdomain in-addr of arpa is used: 14.32.38.138.in-addr.arpa
- A (recursive) lookup for the *pointer* (PTR) record for this will return the name mary.bath.ac.uk
- CIDR causes complications

The Domain Name System

Reverse DNS

- Thus there are actually *two* database to maintain: name-to-number and number-to-name
- It is very common for administrators to forget and these two get out of step: a problem as reverse lookup is sometimes used in authentication of connections

The Domain Name System

Other DNS Data

- The DNS can supply more than just IP addresses (“A” records)
- We have seen PTR records
- There are many other kinds of record defined, but few are used

The Domain Name System

Other DNS Data

Name	Type	RFC	
A	1	1035	IP address
AAAA	28	2874	IPv6 address (128 bit)
NS	2	1035	Authoritative name server
CNAME	5	1035	Canonical name
SOA	6	1035	Start of Authority
PTR	12	1035	Pointer
HINFO	13	1035	Host info
MX	15	1035	Email exchange
SRV	33	2782	Server selection
AXFR	252	1035	Zone transfer
ANY	255	1035	All records

The Domain Name System

Other DNS Data

- HINFO is information about the machine, e.g., operating system. Not often used as this kind of information might possibly be used to aid an attack on the machine
- NS gives an authoritative name server for the domain
- AXFR is for a zone transfer, usually for a primary to a secondary
- ALL is for all records related to a name

The Domain Name System

Other DNS Data

- MX gives the IP address of one or more machines that will accept email for the given domain

The Domain Name System

Other DNS Data

```
% dig bath.ac.uk mx
...
bath.ac.uk.          102931  IN      MX      20 kelly.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      20 roche.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      10 bucs.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      20 binda.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      20 coppi.bath.ac.uk.
```

- The integer before the FQDN is a *preference*: one of the hosts with the smallest preference is contacted first. If that is down, try the next

The Domain Name System

Other DNS Data

```
% dig bath.ac.uk mx
...
bath.ac.uk.          102931  IN      MX      20 kelly.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      20 roche.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      10 bucs.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      20 binda.bath.ac.uk.
bath.ac.uk.          102931  IN      MX      20 coppi.bath.ac.uk.
```

- Note that no machine of the name bath.ac.uk actually exists, but we can still send to somebody@bath.ac.uk

The Domain Name System

Other DNS Data

- SRV is slightly different. Its purpose is *service discovery*, e.g., finding a printer on the local network
- Sending a SRV request for _lpr._tcp.example.com could return the addresses of a machine or machines that are willing to provide a print service
- _www._tcp.example.com for a web server, and so on

The Domain Name System

Other DNS Data

- SRV is not currently well supported, but the Zeroconf Working Group of the IETF is set to use it for *zero configuration networking*
- The idea is to plug in a computer (or general appliance) and have it discover services (printer, web, etc.) for itself with no user configuration

The Domain Name System

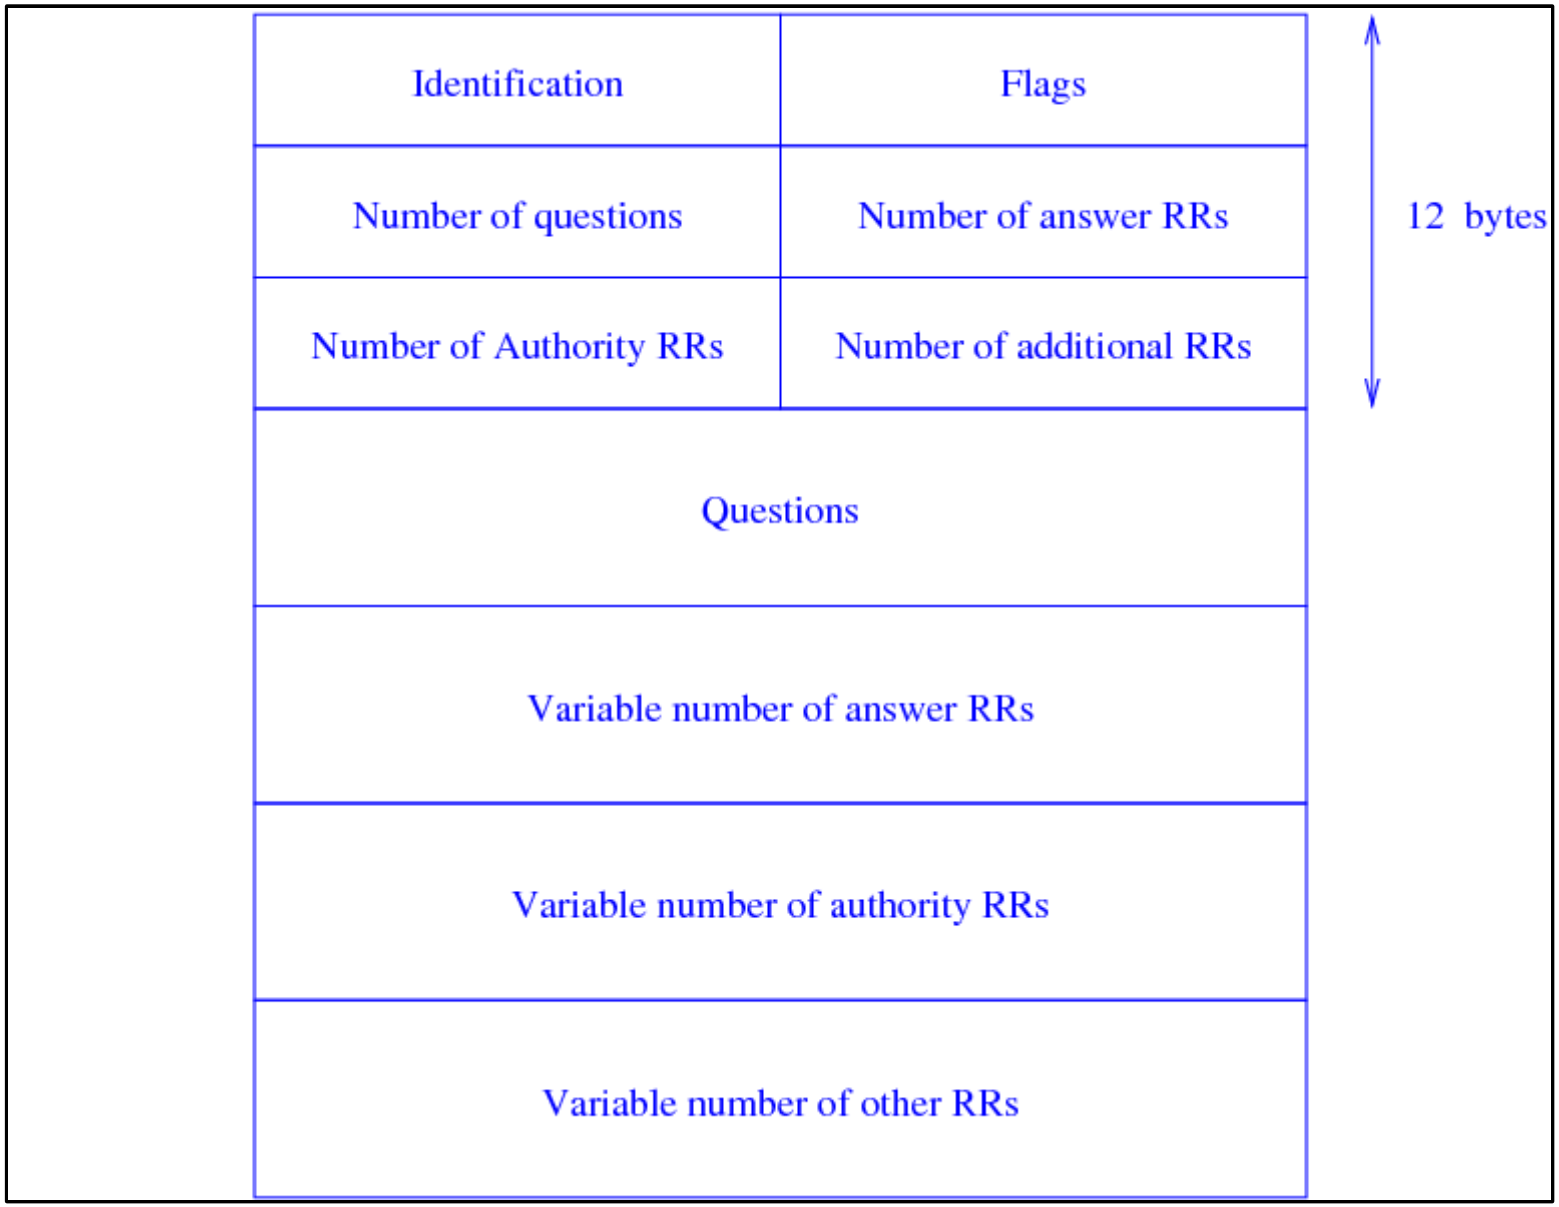

The Domain Name System

DNS Packet Format

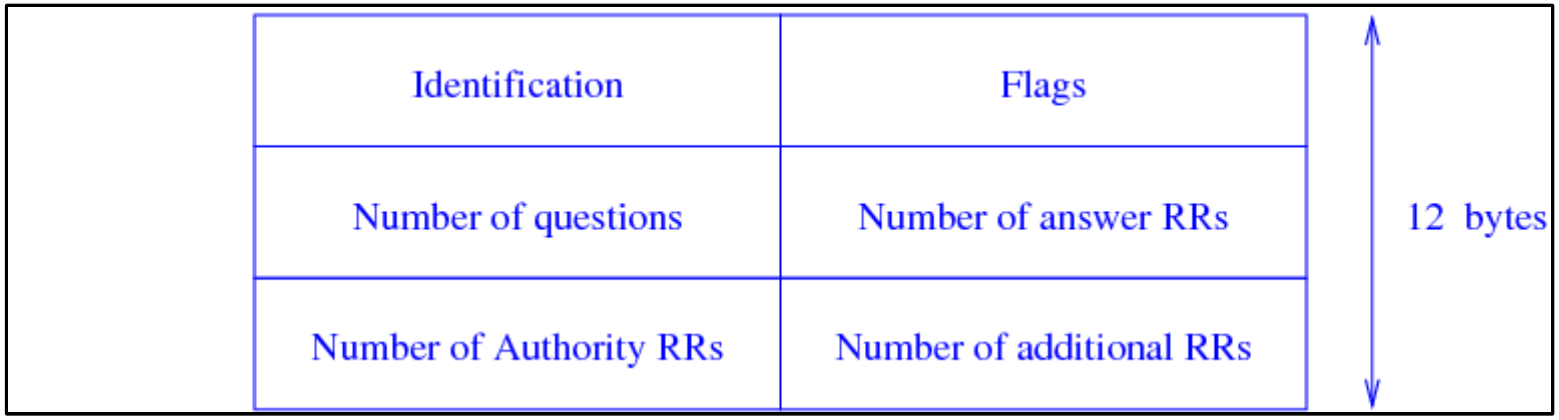

- 12 bytes of header, followed by four variable length fields
- Identification is a value chosen by the client in the query message and returned by the server in the reply. This is used to match replies with questions

The Domain Name System

DNS Packet Format

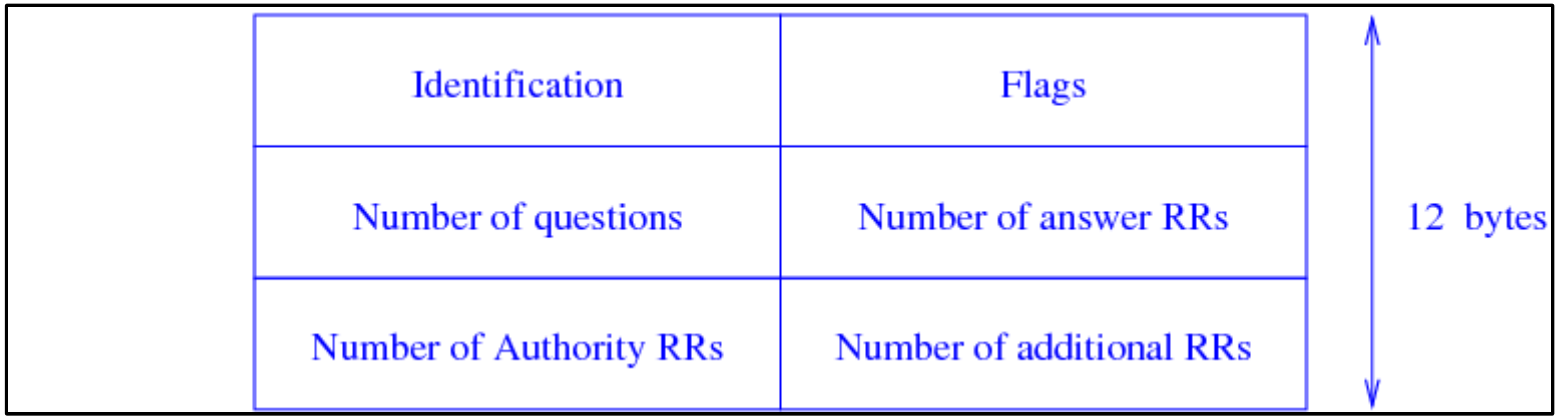

- Flags. A set of bits:
  - QR: 1 for query, 0 for response
  - Opcode: usually 0 for a standard query
  - AA: set on a authoritative answer

The Domain Name System

DNS Packet Format

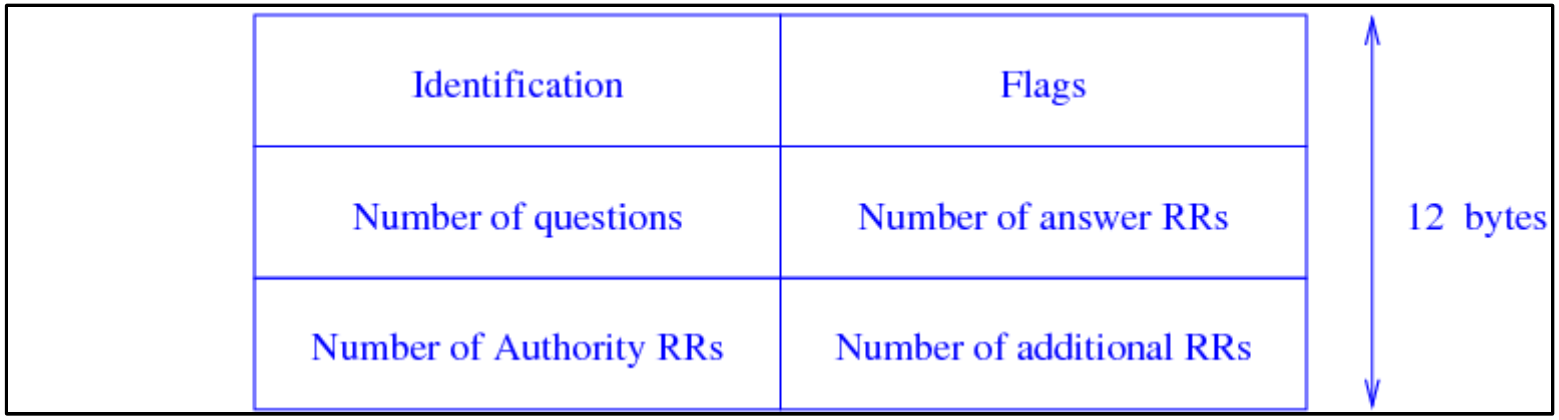

- Flags. A set of bits:
  - TC: truncated. Couldn't fit the answer in 512 bytes (see later)
  - RC: recursion desired. The name server should do a recursive lookup. Usually set

The Domain Name System

DNS Packet Format

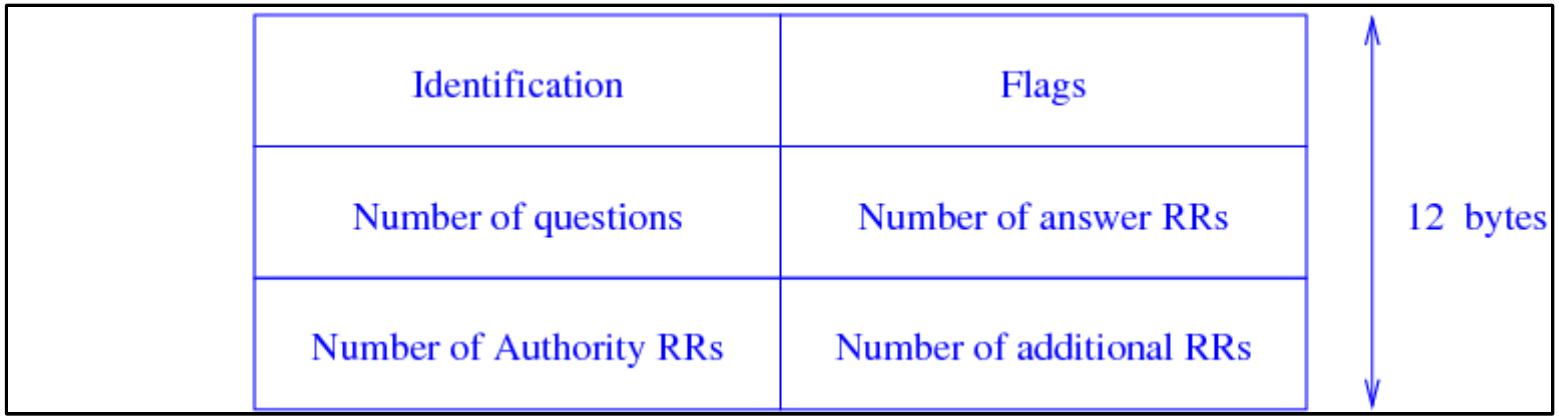

- Flags. A set of bits:
  - RA: recursion available. The server is able to do a recursive lookup. Usually set
  - RC: four bits of return code. 0 is no error, while 3 is *name error*: a response from an authoritative server saying no such name exists

The Domain Name System

DNS Packet Format

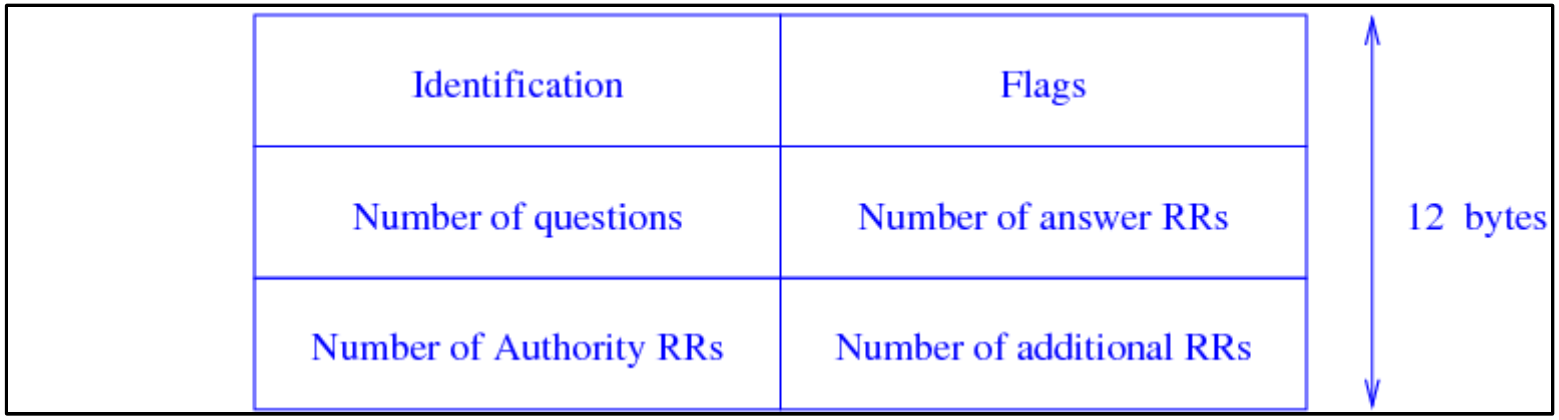

- Next four fields give the number of resource records that follow. This is usually 1,0,0,0 for a request

The Domain Name System

DNS Packet Format

- The question starts with the name we wish to resolve as a sequence of labels, where a label is sent as a single byte containing the length of the label followed by the label, terminated by a 0 byte:

4mary4bath2ac2uk0

The Domain Name System

DNS Question Format

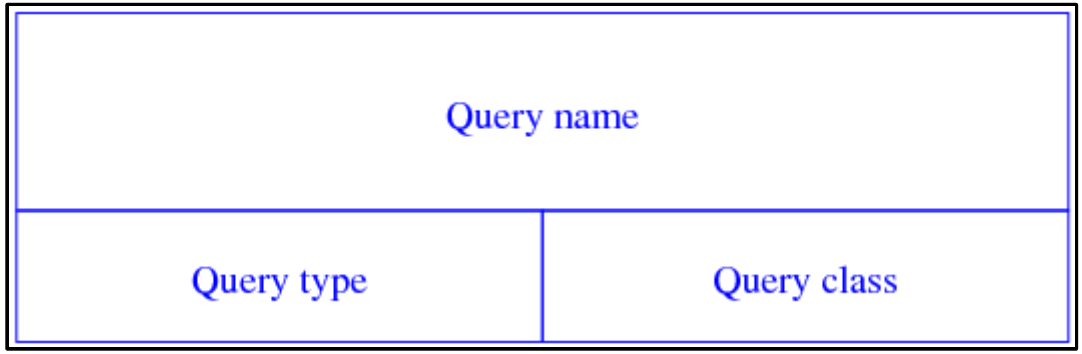

- Query type: a number that specifies A or AAAA or NS, etc
- Query class: normally 1, denoting an IP address

The Domain Name System

DNS Packet Format

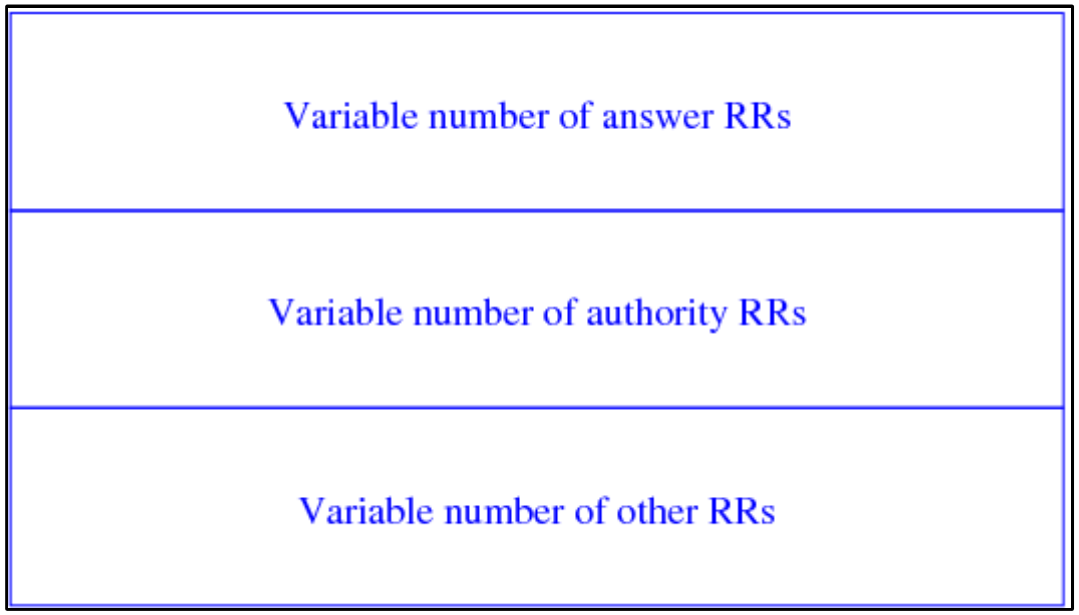

The Domain Name System

DNS Resource Record Format

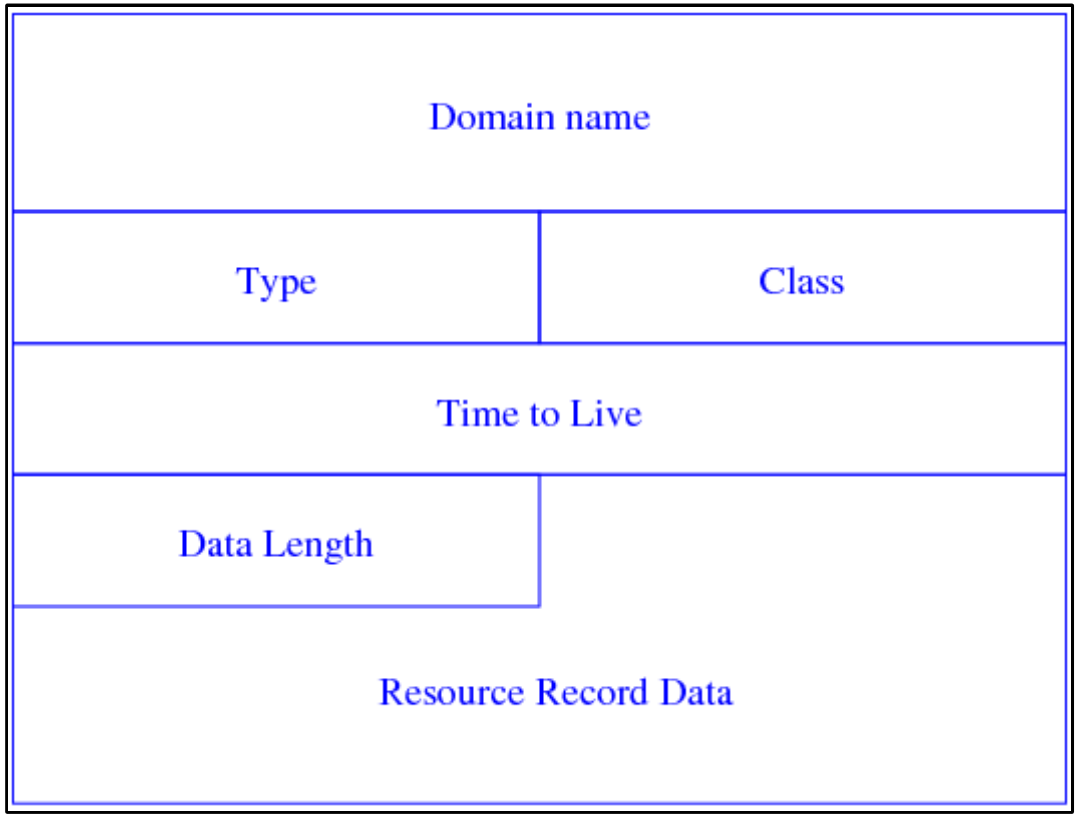

- Domain name, type and class as before

The Domain Name System

DNS Resource Record Format

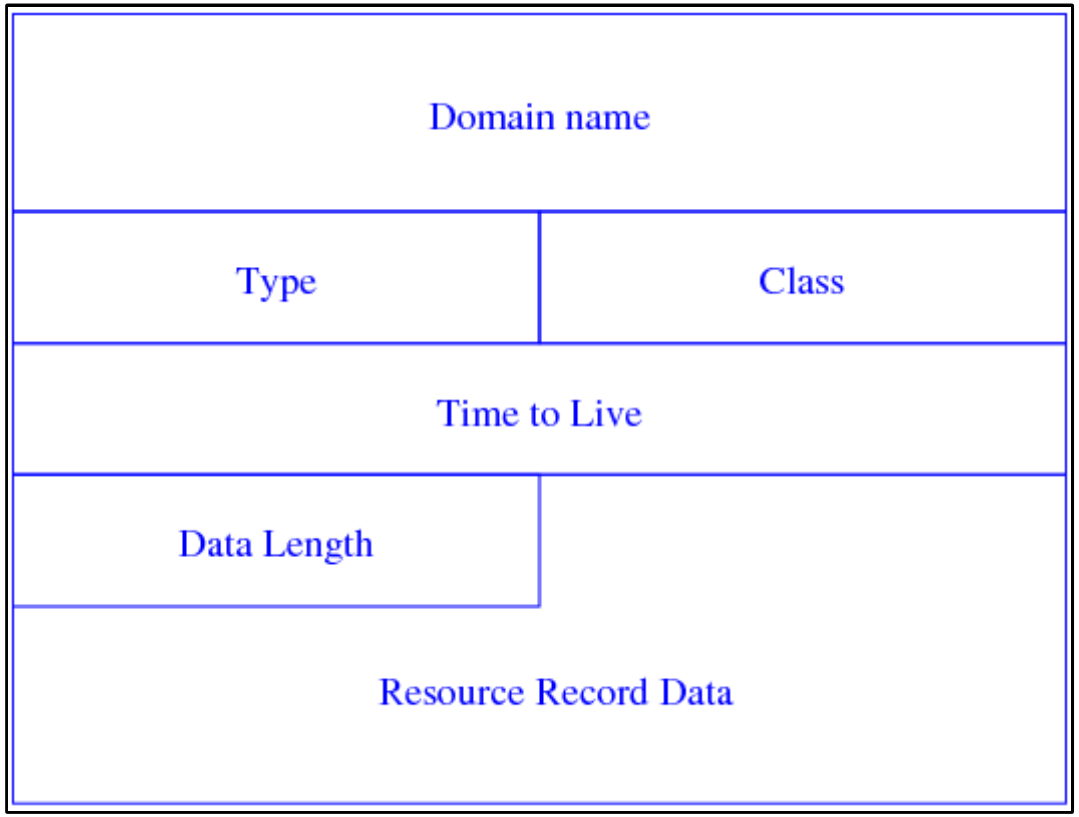

- TTL: in seconds. How long this data should be cached

The Domain Name System

DNS Resource Record Format

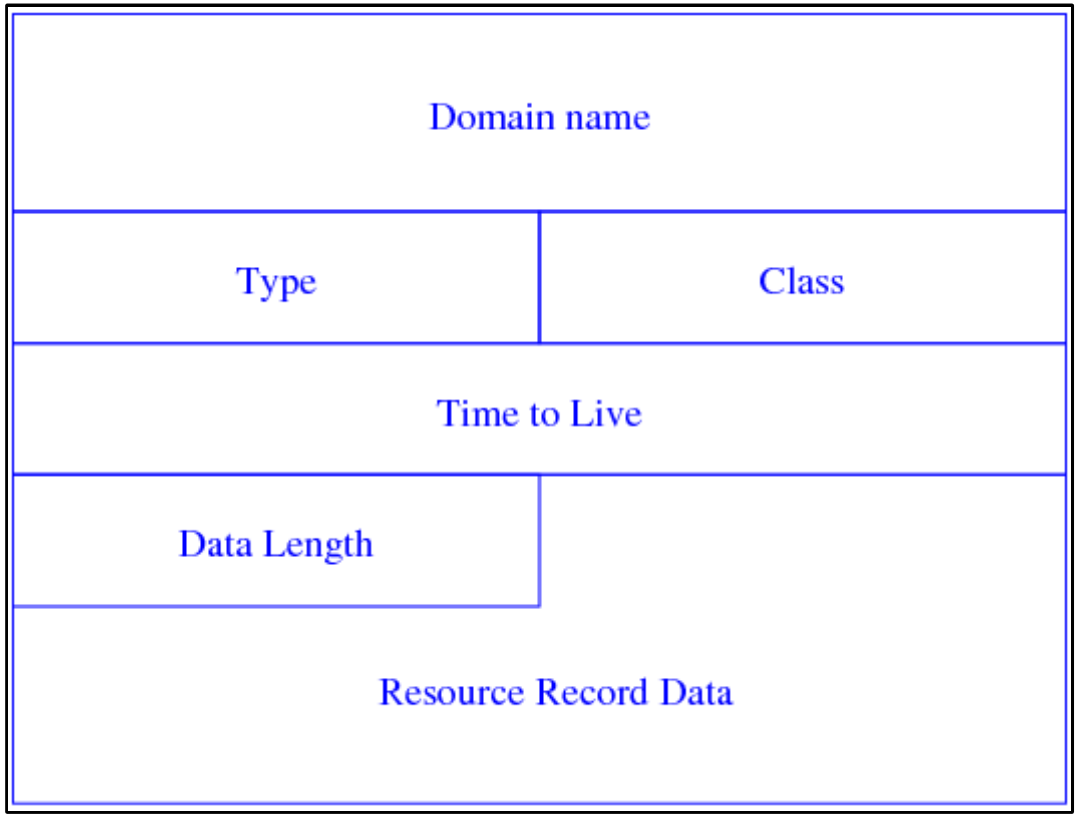

- Resource data is a length followed by the data

The Domain Name System

DNS Resource Record Format

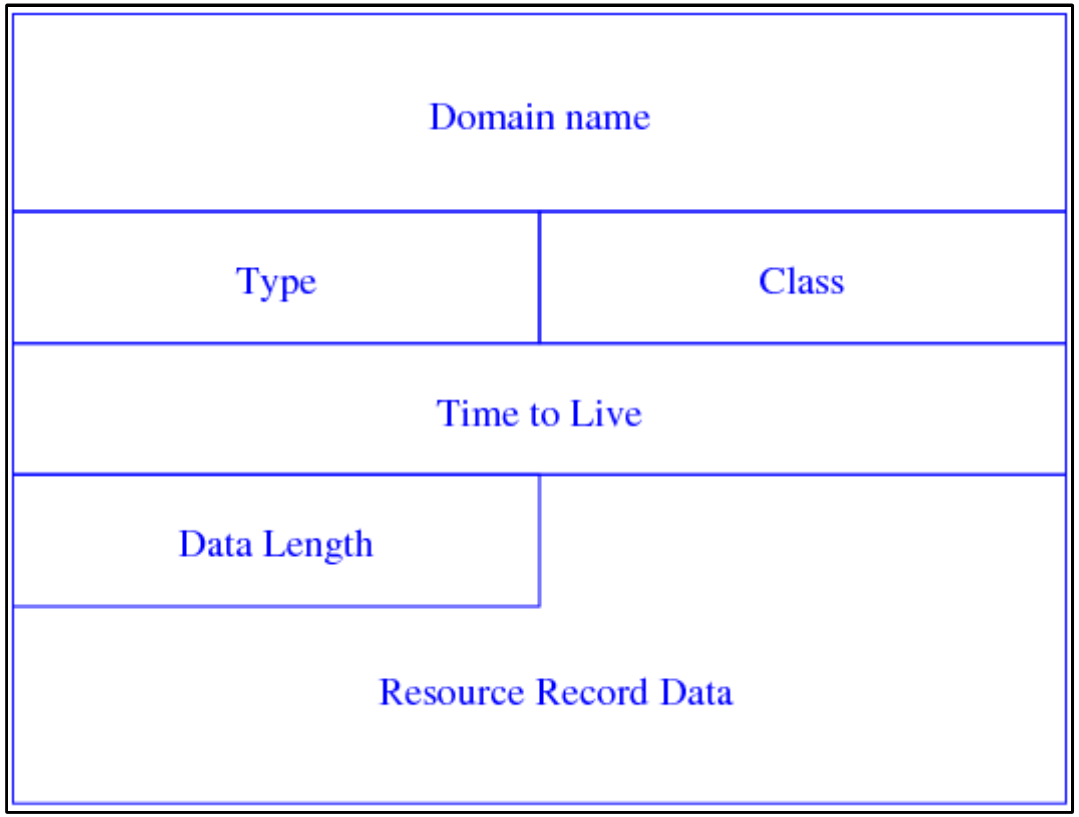

- For an A reply the data is just four bytes of IP address

The Domain Name System

DNS Resource Record Format

- A simple form of compression can be used if the data contains lots of repetition

The Domain Name System

DNS Problems

- There is no authentication: if I get a reply saying that 138.38.32.14 is the address for www.bath.ac.uk can I trust this?
- Some host on the lookup path might have been subverted and made to hand out the wrong address
- Then I would be redirected to someone else's web page still in the belief this was the page I wanted: a problem if, say, I was looking for the page of a bank or shop

The Domain Name System

DNS Problems

- A solution exists in *secure DNS*, which uses cryptography to secure DNS lookups
- Not much in use, yet

The Domain Name System

DNS and Spam

- DNS has been coopted in the fight against spam email: a *Realtime Blackhole List* (RBL) is a database of IP addresses associated with spammers
- The *Mail Abuse Prevention System* (MAPS) uses DNS to look up this database

The Domain Name System

DNS and Spam

- If an email arrives from 10.0.0.1 the receiving host does a DNS lookup of 1.0.0.10.bl.spamcop.net
- Spamcop is a provider of RBL services: others exist
- If the address is in the list, an IP address such as 172.0.0.2 is returned
- If not, “name not found” is returned

The Domain Name System

DNS and Spam

- This allows the host quickly to check if the email is likely from a spammer