Exercise 2.1: Debugging nameservers using dig +norec

You do NOT need to be root to run this exercise. NOTE: it is very good
practice to put a trailing dot after every hostname - this prevents the
default domain from `/etc/resolv.conf` being appended.

This example: testing __www.tiscali.co.uk.__

1. Make a query starting at a root nameserver

The root servers are called `[a-m].root-servers.net.` - pick any one to
start.

```
    $ dig +norec @a.root-servers.net. www.tiscali.co.uk. a

    ; <<>> DiG 9.3.1 <<>> +norec @a.root-servers.net.
www.tiscali.co.uk. a
    ; (1 server found)
    ;; global options:  printcmd
    ;; Got answer:
    ;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 5252
    ;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 9, ADDITIONAL: 10

    ;; QUESTION SECTION:
    ;www.tiscali.co.uk.             IN      A

    ;; AUTHORITY SECTION:
    uk.                     172800  IN      NS      NS1.NIC.uk.
    uk.                     172800  IN      NS      NS2.NIC.uk.
    uk.                     172800  IN      NS      NS4.NIC.uk.
    uk.                     172800  IN      NS      NS3.NIC.uk.
    uk.                     172800  IN      NS      NS5.NIC.uk.
    uk.                     172800  IN      NS      NSA.NIC.uk.
    uk.                     172800  IN      NS      NSB.NIC.uk.
    uk.                     172800  IN      NS      NS6.NIC.uk.
    uk.                     172800  IN      NS      NS7.NIC.uk.

    ;; ADDITIONAL SECTION:
    NS1.NIC.uk.             172800  IN      A       195.66.240.130
    NS2.NIC.uk.             172800  IN      A       217.79.164.131
    NS4.NIC.uk.             172800  IN      AAAA
2001:630:181:35::83
    NS4.NIC.uk.             172800  IN      A       194.83.244.131
    NS3.NIC.uk.             172800  IN      A       213.219.13.131
    NS5.NIC.uk.             172800  IN      A       213.246.167.131
    NSA.NIC.uk.             172800  IN      A       204.74.112.44
    NSB.NIC.uk.             172800  IN      A       204.74.113.44
    NS6.NIC.uk.             172800  IN      A       213.248.254.130
    NS7.NIC.uk.             172800  IN      A       212.121.40.130
```

```
    ;; Query time: 98 msec
    ;; SERVER: 198.41.0.4#53(198.41.0.4)
    ;; WHEN: Mon May 16 12:34:06 2005
    ;; MSG SIZE  rcvd: 373
```

Note: We only got back NS records (plus some related information - the A
records which correspond to those nameservers). This is a REFERRAL.

In theory we should repeat this query for `b.root-servers.net`, `c.root-servers.net` ... and check we get the same answers. Occasionally
you _might_ find inconsistencies between root servers, but it's rare.

2. Note the nine nameservers we saw in the response

(Remember that DNS names are not case sensitive. We also get them back in a
random order; this doesn't matter because we are going to try every one
anyway)

```
>     ns1.nic.uk.
>     ns2.nic.uk.
>     ns3.nic.uk.
>     ns4.nic.uk.
>     ns5.nic.uk.
>     ns6.nic.uk.
>     ns7.nic.uk.
>     nsA.nic.uk.
>     nsB.nic.uk.
```

3. Repeat the query for all NS records in turn

```
    $ dig +norec @ns1.nic.uk. www.tiscali.co.uk. a

    ; <<>> DiG 9.3.1 <<>> +norec @ns1.nic.uk. www.tiscali.co.uk. a
    ; (1 server found)
    ;; global options:  printcmd
    ;; Got answer:
    ;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 28452
    ;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 2, ADDITIONAL: 1

    ;; QUESTION SECTION:
    ;www.tiscali.co.uk.             IN      A

    ;; AUTHORITY SECTION:
    tiscali.co.uk.          172800  IN      NS      ns0.as9105.com.
    tiscali.co.uk.          172800  IN      NS
ns0.tiscali.co.uk.
```

```
    ;; ADDITIONAL SECTION:
    ns0.tiscali.co.uk.      172800  IN      A       212.74.114.132

    ;; Query time: 20 msec
    ;; SERVER: 195.66.240.130#53(195.66.240.130)
    ;; WHEN: Mon May 16 12:37:23 2005
    ;; MSG SIZE  rcvd: 97


    $ dig +norec @ns2.nic.uk. www.tiscali.co.uk. a
    ... results snipped to save paper

    $ dig +norec @ns3.nic.uk. www.tiscali.co.uk. a
    ... results snipped to save paper
    ... etc
```

Check the results are consistent!

Note: if a server is authoritative for both a domain and a subdomain, it
will immediately return the result for the subdomain. This is OK. In this
example, the same servers are authoritative for both `.uk` and `.co.uk`,
so they can delegate us immediately to the servers for `tiscali.co.uk`, taking
us down two levels of the DNS hierarchy in one go.

You can see here that we are getting another delegation, this time to two
other nameservers:

```
>     ns0.as9105.com
>     ns0.tiscali.co.uk
```

4. Continue to repeat the query for all NS records found in step 3

```
    $ dig +norec @ns0.tiscali.co.uk. www.tiscali.co.uk. a

    ; <<>> DiG 9.3.1 <<>> +norec @ns0.tiscali.co.uk.
www.tiscali.co.uk. a
    ; (1 server found)
    ;; global options:  printcmd
    ;; Got answer:
    ;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 35827
    ;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL:
2

    ;; QUESTION SECTION:
    ;www.tiscali.co.uk.             IN      A

    ;; ANSWER SECTION:
```

```
    www.tiscali.co.uk.      3600    IN      A       212.74.101.10

    ;; AUTHORITY SECTION:
    tiscali.co.uk.          3600    IN      NS      ns0.as9105.com.
    tiscali.co.uk.          3600    IN      NS
ns0.tiscali.co.uk.

    ;; ADDITIONAL SECTION:
    ns0.as9105.com.         2419200 IN      A       212.139.129.130
    ns0.tiscali.co.uk.      2419200 IN      A       212.74.114.132

    ;; Query time: 21 msec
    ;; SERVER: 212.74.114.132#53(212.74.114.132)
    ;; WHEN: Mon May 16 12:40:00 2005
    ;; MSG SIZE  rcvd: 129


    $ dig +norec @ns0.as9105.com. www.tiscali.co.uk. a
    ...
    ;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL:
2
    ...
    ;; ANSWER SECTION:
    www.tiscali.co.uk.      3600    IN      A       212.74.101.10
```

This time, instead of getting another delegation, we have found the answer
we are looking for. Note that the nameservers are both giving authoritative
answers (`flags: aa`), and the results are the same. Also note that the
'AUTHORITY SECTION' in the response has the *same* list of nameservers as we
used to perform the query. (This second set of NS records are contained
within the authoritative server itself, as opposed to the delegation from
above)

5. Checklist

* Were all the nameservers reachable?
* Were there at least two nameservers on two different subnets?
* Did they all give either a referral or an AA (Authoritative Answer)?
* Were all the answers the same?
* Were the TTL values reasonable?
* Does the final list of nameservers in the AUTHORITY SECTION match the
    list of nameservers in the referral?

6. Now check the NS records themselves!

Notice that every NS record points to the NAME of a host, not an IP address. (It is illegal for an NS record to point at an IP address, it will
not work at all)

However, when we issued a command like `dig @ns0.as9105.com ...`, we were
relying on dig converting this name to the correct IP address. It performs a
recursive lookup to find the IP address of this server, so that it can send
the query there.

Therefore, you need to start again and check every NS record you found,
starting from the root again, in exactly the same way! This is tedious, and
usually the top-level servers are right. But it's worth checking your
country-level NS records and your own NS records.

Example: check ns0.as9105.com

```
    $ dig +norec @a.root-servers.net. ns0.as9105.com. a
    ... referral to [a-m].gtld-servers.net.

    $ dig +norec @a.gtld-servers.net. ns0.as9105.com. a
    ;; flags: qr; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 1
    ;; ANSWER SECTION:
    ns0.as9105.com.         172800  IN      A       212.139.129.130
<====

    ;; AUTHORITY SECTION:
    as9105.com.             172800  IN      NS      ns0.as9105.com.
    as9105.com.             172800  IN      NS
ns0.tiscali.co.uk.
```

Notice that here we got an answer - but it is not an authoritative answer!
(As well as 'aa' missing, notice that the machine we queried is not one of
the machines listed in the 'authority section')

This is not an error as long as the answer is correct - it's called a "glue
record" which we'll discuss later - but we need to continue downwards to
find the true authoritative source:

```
    $ dig +norec @ns0.as9105.com. ns0.as9105.com. a
    ;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL:
1
```

```
    ;; ANSWER SECTION:
    ns0.as9105.com.         2419200 IN      A       212.139.129.130
<====

    ;; AUTHORITY SECTION:
    as9105.com.             600     IN      NS
ns0.tiscali.co.uk.
    as9105.com.             600     IN      NS      ns0.as9105.com.

    ;; ADDITIONAL SECTION:
    ns0.tiscali.co.uk.      2419200 IN      A       212.74.114.132


    $ dig +norec @ns0.tiscali.co.uk. ns0.as9105.com. a
    ;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL:
1

    ;; ANSWER SECTION:
    ns0.as9105.com.         2419200 IN      A       212.139.129.130
<====

    ;; AUTHORITY SECTION:
    as9105.com.             600     IN      NS
ns0.tiscali.co.uk.
    as9105.com.             600     IN      NS      ns0.as9105.com.

    ;; ADDITIONAL SECTION:
    ns0.tiscali.co.uk.      2419200 IN      A       212.74.114.132
```

Now we check:

* Were all the answers the same? (Yes: 212.139.129.130 from both `a.gtld-servers.net` and the authoritative nameservers)
* Did the delegation match the NS records in the authoritative nameservers? (Yes: delegation to `ns0.as9105.com` and `ns0.tiscali.co.uk`, and these records were also given in the 'authority section' of the final response)

Negative answers

The non-existence of a RR is an important piece of information too. The
response you get should look like this:

```
    $ dig +norec @ns0.tiscali.co.uk. wibble.tiscali.co.uk. a
    ;; ->>HEADER<<- opcode: QUERY, status: NXDOMAIN, id: 51165
    ;; flags: qr aa; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL:
0

    ;; AUTHORITY SECTION:
    tiscali.co.uk.          3600    IN      SOA
ns0.tiscali.co.uk.
    hostmaster.uk.tiscali.com. 2005051301 10800 3600 604800 3600
```

AA is set, but there is nothing in the answer apart from the SOA. The
parameters in the SOA are used to work out how much negative caching is
allowed. (Old caches use the TTL of the SOA itself; new caches uses the SOA
'minimum' value. It's best to set both to the same value. We'll look at the
exact format of the SOA record shortly.)

Meaning of flags (from RFC 1034/RFC 1035)

QR	A one bit field that specifies whether this message is a query (0), or a response (1).
AA	Authoritative Answer - this bit is valid in responses, and specifies that the responding name server is an authority for the domain name in question section.
RD	Recursion Desired - this bit may be set in a query and is copied into the response. If RD is set, it directs the name server to pursue the query recursively. Recursive query support is optional.
RA	Recursion Available - this be is set or cleared in a response, and denotes whether recursive query support is available in the name server.

As well as the lack of 'AA' flag, a good way to spot cached answers is to repeat the query a few times and watch the TTL counting downwards.

```
$ dig psg.com.
;; ANSWER SECTION:
psg.com.                14397   IN      A       147.28.0.62

$ dig psg.com.
;; ANSWER SECTION:
psg.com.                14384   IN      A       147.28.0.62
```