

DNS lab: dig, part 2

Debugging nameservers using dig +nored

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.__`

For this lab, we'll need to temporarily change your default nameserver, configured in `/etc/resolv.conf`, to `10.20.0.254`, like so:

```
# ee /etc/resolv.conf
```

... and set the nameserver to be:

```
nameserver 10.20.0.254
```

Save the file and exit the editor.

Note: we need to do this, otherwise we won't be able to lookup names on the Internet!

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 +nored @f.root-servers.net www.tiscali.co.uk. a

; <<>> DiG 9.7.2-P3 <<>> +nored @a.root-servers.net. www.tiscali.co.uk. a
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 8712
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 11, ADDITIONAL: 14

;; 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  ns3.nic.uk.
uk.      172800  IN  NS  ns4.nic.uk.
uk.      172800  IN  NS  ns5.nic.uk.
uk.      172800  IN  NS  ns6.nic.uk.
uk.      172800  IN  NS  ns7.nic.uk.
uk.      172800  IN  NS  nsa.nic.uk.
uk.      172800  IN  NS  nsb.nic.uk.
uk.      172800  IN  NS  nsc.nic.uk.
uk.      172800  IN  NS  nsd.nic.uk.

;; ADDITIONAL SECTION:
ns1.nic.uk.      172800  IN  AAAA  2a01:40:1001:35::2
ns1.nic.uk.      172800  IN  A     195.66.240.130
ns2.nic.uk.      172800  IN  A     217.79.164.131
ns3.nic.uk.      172800  IN  A     213.219.13.131
ns4.nic.uk.      172800  IN  AAAA  2001:630:181:35::83
```

```

ns4.nic.uk.      172800 IN A   194.83.244.131
ns5.nic.uk.      172800 IN A   213.246.167.131
ns6.nic.uk.      172800 IN A   213.248.254.130
ns7.nic.uk.      172800 IN A   212.121.40.130
nsa.nic.uk.      172800 IN AAAA 2001:502:ad09::3
nsa.nic.uk.      172800 IN A   156.154.100.3
nsb.nic.uk.      172800 IN A   156.154.101.3
nsc.nic.uk.      172800 IN A   156.154.102.3
nsd.nic.uk.      172800 IN A   156.154.103.3

```

```

;; Query time: 8 msec
;; SERVER: 198.41.0.4#53(198.41.0.4)
;; WHEN: Tue Feb 15 15:53:13 2011
;; MSG SIZE rcvd: 497

```

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 eleven 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.
nsc.nic.uk.
nsd.nic.uk.

```

3. Repeat the query for all NS records in turn

```

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

; <<>> DiG 9.7.2-P3 <<>> +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
...
```

```
$ dig +norec @ns3.nic.uk. www.tiscali.co.uk. a
...
... etc
```

Check that 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 refer 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 9.7.2-P3 <<>> +norec @ns0.tiscali.co.uk. www.tiscali.co.uk. a
; (1 server found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 52841
;; 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.      300 IN    A      212.74.99.30

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

;; ADDITIONAL SECTION:
ns0.as9105.com.         604800 IN     A      212.139.129.130
ns0.tiscali.co.uk.     604800 IN     A      212.74.114.132

;; Query time: 322 msec
```

```
;; SERVER: 212.74.114.132#53(212.74.114.132)
;; WHEN: Tue Feb 15 16:01:04 2011
;; MSG SIZE rcvd: 129
```

```
$ dig +noredc @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. 300 IN A 212.74.99.30
```

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)

Hint: try this!

```
$ dig +nssearch tiscali.co.uk
```

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. In fact, we are doing two queries:

- dig asks for the IP address of ns0.as9105.com, performing a recursive lookup using the nameserver listed in /etc/resolv.conf
- once dig has gotten the IP address of the nameserver, dig can send its query to that server

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

```
$ dig +noredc @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 explain later - but we need to continue downwards to
find the true authoritative source:

```
$ dig +noredc @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 +noredc @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)

The meaning of NOERROR

You may have paid attention to the status: field of the dig output:

```
status: NXDOMAIN
or
status: NOERROR
```

NXDOMAIN means Non-existent Domain - it means: "Sorry, no data exists for the given NAME at all". It basically means that there is no DNS data for the name you're querying. For instance:

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

... will return NXDOMAIN. There is nothing at all for "wibble" under tiscali.co.uk. No A record, no AAAA record, etc...

Now, you may also have noticed that the ANSWER section can contain 0 answers, but still, the queried server returns NOERROR, and not NXDOMAIN.

Why is this ?

Let's say for example that we want to know the IP address for www.tiscali.co.uk:

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

So far so good - you should see:

```
status: NOERROR
ANSWER: 1
```

Now, let's ask for a *different* type ("Resource Record Type, formally speaking):

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

Notice that we ask for a TXT record for the name "www.tiscali.co.uk."

What do we get ?

```
status: NOERROR
ANSWER: 0
```

How can this be ?

NOERROR in this case means "Sorry, no data exists for the given NAME & TYPE requested". Aha! Here we're being told that there is no TXT record data for www.tiscali.co.uk - but there may be other data under other data types.

Indeed, we know from earlier that:

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

... will return us the IP address of www.tiscali.co.uk.

Therefore, a non-existent name (NXDOMAIN) or an empty answer (NOERROR, ANSWER: 0) is *still* an answer, and we need to remember (cache) this, as we'll see below.

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.talktalkplc.com.
2011012703 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.

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 bit 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
                ^^^^^
```

Other dig options

Other dig options you may want to try - use the manpage to find out what they do!

dig +tcp

```
dig +trace
```

Try other options you find in the man page!

Clean up

Finally, when you're done, remember to restore your `/etc/resolv.conf`:

```
nameserver 10.20.0.230
```