

# 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

```
$ dig +norec @a.root-servers.net. www.tiscali.co.uk. a
; <<>> DiG 8.3 <<>> +norec @a.root-servers.net. www.tiscali.co.uk. a
; (1 server found)
;; res options: init defnam dnsrch
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 29971
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 5, ADDITIONAL: 6
;; QUERY SECTION:
;;      www.tiscali.co.uk, type = A, class = IN

;; AUTHORITY SECTION:
uk.          2D IN NS      NS1.NIC.uk.
uk.          2D IN NS      NS0.JA.NET.
uk.          2D IN NS      NS.UU.NET.
uk.          2D IN NS      SEC-NOM.DNS.UK.PSI.NET.
uk.          2D IN NS      NS2.NIC.uk.

;; ADDITIONAL SECTION:
NS1.NIC.uk.  2D IN A      195.66.240.130
NS0.JA.NET.  2D IN A      128.86.1.20
NS0.JA.NET.  2D IN A      193.63.94.20
NS.UU.NET.   2D IN A      137.39.1.3
SEC-NOM.DNS.UK.PSI.NET. 2D IN A      154.32.105.90
NS2.NIC.uk.  2D IN A      217.79.164.131

;; Total query time: 662 msec
;; FROM: vaio.linnet.org to SERVER: a.root-servers.net. 198.41.0.4
;; WHEN: Mon Jun  9 09:31:00 2003
;; MSG SIZE sent: 35 rcvd: 248
```

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 etc and check we get the same answers. Occasionally root servers *do* have problems.

## 2. Note the five nameservers we saw in the response

(Remember that DNS names are not case sensitive)

```
ns1.nic.uk.
ns0.ja.net.
ns.uu.net.
sec-nom.dns.uk.psi.net.
ns2.nic.uk.
```

## 3. Repeat the query for all NS records in turn

```

$ dig +norec @ns1.nic.uk. www.tiscali.co.uk. a
; <<>> DiG 8.3 <<>> +norec @ns1.nic.uk. www.tiscali.co.uk. a
; (1 server found)
;; res options: init defnam dnsrch
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 15102
;; flags: qr; QUERY: 1, ANSWER: 0, AUTHORITY: 2, ADDITIONAL: 1
;; QUERY SECTION:
;;      www.tiscali.co.uk, type = A, class = IN

;; AUTHORITY SECTION:
tiscali.co.uk.      2D IN NS      ns0.tiscali.co.uk.
tiscali.co.uk.      2D IN NS      ns0.as9105.com.

;; ADDITIONAL SECTION:
ns0.tiscali.co.uk.  2D IN A      212.74.114.132  <-- "Glue record"

;; Total query time: 757 msec
;; FROM: vaio.linnet.org to SERVER: ns1.nic.uk. 195.66.240.130
;; WHEN: Mon Jun  9 09:31:25 2003
;; MSG SIZE  sent: 35  rcvd: 97

$ dig +norec @ns0.ja.net. www.tiscali.co.uk. a
... results snipped to save paper
$ dig +norec @ns.uu.net. www.tiscali.co.uk. a
... results
$ dig +norec @sec-nom.dns.uk.psi.net. www.tiscali.co.uk. a
... results
$ dig +norec @ns2.nic.uk. www.tiscali.co.uk. a
... results

```

**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)

#### 4. Continue to repeat the query for all NS records found

```

$ dig +norec @ns0.tiscali.co.uk. www.tiscali.co.uk. a
;; got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 57638
;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 2
;; QUERY SECTION:
;;      www.tiscali.co.uk, type = A, class = IN

;; ANSWER SECTION:
www.tiscali.co.uk.  1H IN A      212.74.101.10

;; AUTHORITY SECTION:
tiscali.co.uk.      6H IN NS      ns0.as9105.com.
tiscali.co.uk.      6H IN NS      ns0.tiscali.co.uk.

;; ADDITIONAL SECTION:
ns0.as9105.com.     1D IN A      212.139.129.130
ns0.tiscali.co.uk.  6H IN A      212.74.114.132

$ dig +norec @ns0.as9105.com. www.tiscali.co.uk. a
...
;; ANSWER SECTION:
www.tiscali.co.uk.  1H IN A      212.74.101.10
... check it's the same

```

We have now found the answer. Also check that the 'AUTHORITY SECTION' in the response has the **same** list of nameservers as we used to perform the query. (These are the NS records contained within the authoritative server itself)

## 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 during each 'dig', we were relying on dig converting (say) 'ns0.as9105.com' 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, in exactly the same way! This is tedious, and generally the top-level servers are right. But it's worth checking your country-level NS records and your own NS records.

### Example: 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.          2D IN A           212.139.129.130

;; AUTHORITY SECTION:
as9105.com.             2D IN NS          ns0.as9105.com.
as9105.com.             2D IN NS          ns0.tiscali.co.uk.
```

But this 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, but we need to continue downwards to find the true authoritative source:

```
$ dig +norec @ns0.as9105.com. as9105.com. a
;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 1
;; ANSWER SECTION:
ns0.as9105.com.          1D IN A           212.139.129.130

;; AUTHORITY SECTION:
as9105.com.             1D IN NS          ns0.as9105.com.
as9105.com.             1D IN NS          ns0.tiscali.co.uk.

$ dig +norec @ns0.tiscali.co.uk. as9105.com. a
;; flags: qr aa; QUERY: 1, ANSWER: 1, AUTHORITY: 2, ADDITIONAL: 1
;; ANSWER SECTION:
ns0.as9105.com.          1D IN A           212.139.129.130
```

```
;; AUTHORITY SECTION:
as9105.com.          1D IN NS      ns0.tiscali.co.uk.
as9105.com.          1D IN NS      ns0.as9105.com.
```

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 listed in the zone)

## 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: 4
;; flags: qr aa; QUERY: 1, ANSWER: 0, AUTHORITY: 1, ADDITIONAL: 0
;; QUERY SECTION:
;;      wibble.tiscali.co.uk, type = A, class = IN

;; AUTHORITY SECTION:
tiscali.co.uk.      1H IN SOA      ns0.tiscali.co.uk. hostmaster.uk.tiscali
                    2003060301    ; serial
                    3H      ; refresh
                    1H      ; retry
                    1W      ; expiry
                    1H )    ; minimum
```

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)

## 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 @noc.ws.afnog.org. psg.com.
;; ANSWER SECTION:
psg.com.            53m44s IN A      147.28.0.62
```

\$ dig @noc.ws.afnog.org. psg.com.

;; ANSWER SECTION:

psg.com.

**53m38s** IN A

147.28.0.62