

F-Root Incident Report 2020-01

Ray Bellis, Director of DNS Operations, 2020-02-22



Summary

A single report was received by ISC from a large network operator on 2020-01-23 at 17:33 UTC of sporadic failures resolving the .net zone, apparently caused by missing glue records for the .net nameservers in some responses from both E (NASA) and F (ISC) root servers. This meant that when a query was sent for the NS records for the .net Top Level Domain (TLD), that in some cases no A and AAAA resource records for the TLD were returned (see Figure 1 for a correct response and Figure 2 for the sporadic insufficient response).

The report was immediately verified and confirmed to be an issue with the nodes operated through our partnership with Cloudflare. The issue was escalated to them straight away and the diagnosis was confirmed. Cloudflare proceeded to both work on the underlying software fault and withdrawal of the BGP announcements of the F-Root prefixes (192.5.5.0/24 and 2001:500:2f::/48) by AS13335. By 20:51 UTC a software fix had been deployed and the routes re-announced and full service was resumed. The time to restore from incident awareness was 3 hours, 18 minutes.

Timeline

2020-01-21 13:55	Phased release of code to fix a software bug across CF's network commences
2020-01-23 17:33	Email report of missing glue received by ISC. The missing glue is reported to be causing sporadic DNS resolution failures, especially for the .net zone.
2020-01-23 17:41	Email acknowledged
2020-01-23 17:46	Report verified and escalated to CF
2020-01-23 17:53	CF internal team began investigating
2020-01-23 18:07	Root cause identified as a new bug introduced in the release from 2020-01-21
2020-01-23 18:34	BGP prefix withdrawal requested
2020-01-23 18:36	Software update to resolve issue initiated
2020-01-23 19:20	BGP prefix withdrawal completed[*]
2020-01-23 19:55	Issues reported with withdrawal from Luxemburg and NZ pops
2020-01-23 20:11	BGP prefix withdrawal fully completed
2020-01-23 20:34	CF authorized to resume service

2020-01-23 20:42	Phased reintroduction of BGP prefix announcements
2020-01-23 20:51	BGP announcements fully rolled out - service restored
2020-01-23 23:47	Traffic volumes reconverge to normal levels

Observations

Effects Caused

Only one report was received, from an operator believed to be using a commercial resolver implementation. The operator reported, based on information from their vendor, that some other users of that resolver platform may have been seeing similar issues, but this was not conclusively verified. It's also understood that clients were seeing timeouts and would typically then retry the query against a different resolver that would likely be serving correct results. Mitigation was possible by flushing affected zones from the resolvers' caches until glue was received from elsewhere at which point resolution would work until the 2 day TTL on the NS records expired. In working to mitigate the impact of the issue, the operator reported to ISC that they unintentionally made a configuration change error and temporarily exacerbated the issue for their users.

ISC ran tests against BIND to check its behaviour when receiving a response without glue and confirmed that it automatically resends the query to an alternate authoritative server. However, it seems clear that not all resolvers will react in the same manner in this situation and the responses were clearly not RFC-compliant, so the behaviour in this circumstance is undefined.

Lessons Learned

Notwithstanding the limited side-effects of the issue attributable to F-Root, in retrospect we should have initiated the withdrawal of the route prefixes from BGP as soon as it was identified that incomplete / incorrect data was being served. The withdrawal of routes did not go as smoothly as expected and Cloudflare and ISC have agreed to perform regular tests to exercise that function.

While a test suite was already in place, in practise it is nearly impossible to exercise every possible DNS conformance requirement. The test suite has been updated to include tests for missing glue, and ISC and Cloudflare will work to devise further conformance tests.

APPENDIX

Figure 1 - .net TLD NS response includes A and AAAA RRs

```
; <<> DiG 9.10.6 <<> @a.root-servers.net NS net +nsid
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->HEADER<- opcode: QUERY, status: NOERROR, id: 4695
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 13, ADDITIONAL: 27
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
;; QUESTION SECTION:
;net.                IN      NS

;; AUTHORITY SECTION:
net.                172800 IN      NS      a.gtld-servers.net.
net.                172800 IN      NS      b.gtld-servers.net.
net.                172800 IN      NS      c.gtld-servers.net.
net.                172800 IN      NS      d.gtld-servers.net.
net.                172800 IN      NS      e.gtld-servers.net.
net.                172800 IN      NS      f.gtld-servers.net.
net.                172800 IN      NS      g.gtld-servers.net.
net.                172800 IN      NS      h.gtld-servers.net.
net.                172800 IN      NS      i.gtld-servers.net.
net.                172800 IN      NS      j.gtld-servers.net.
net.                172800 IN      NS      k.gtld-servers.net.
net.                172800 IN      NS      l.gtld-servers.net.
net.                172800 IN      NS      m.gtld-servers.net.

;; ADDITIONAL SECTION:
a.gtld-servers.net. 172800 IN      A       192.5.6.30
b.gtld-servers.net. 172800 IN      A       192.33.14.30
c.gtld-servers.net. 172800 IN      A       192.26.92.30
d.gtld-servers.net. 172800 IN      A       192.31.80.30
e.gtld-servers.net. 172800 IN      A       192.12.94.30
f.gtld-servers.net. 172800 IN      A       192.35.51.30
g.gtld-servers.net. 172800 IN      A       192.42.93.30
h.gtld-servers.net. 172800 IN      A       192.54.112.30
i.gtld-servers.net. 172800 IN      A       192.43.172.30
j.gtld-servers.net. 172800 IN      A       192.48.79.30
k.gtld-servers.net. 172800 IN      A       192.52.178.30
l.gtld-servers.net. 172800 IN      A       192.41.162.30
m.gtld-servers.net. 172800 IN      A       192.55.83.30
a.gtld-servers.net. 172800 IN      AAAA    2001:503:a83e::2:30
b.gtld-servers.net. 172800 IN      AAAA    2001:503:231d::2:30
c.gtld-servers.net. 172800 IN      AAAA    2001:503:83eb::30
d.gtld-servers.net. 172800 IN      AAAA    2001:500:856e::30
e.gtld-servers.net. 172800 IN      AAAA    2001:502:1ca1::30
f.gtld-servers.net. 172800 IN      AAAA    2001:503:d414::30
g.gtld-servers.net. 172800 IN      AAAA    2001:503:eea3::30
h.gtld-servers.net. 172800 IN      AAAA    2001:502:8cc::30
i.gtld-servers.net. 172800 IN      AAAA    2001:503:39c1::30
j.gtld-servers.net. 172800 IN      AAAA    2001:502:7094::30
k.gtld-servers.net. 172800 IN      AAAA    2001:503:d2d::30
l.gtld-servers.net. 172800 IN      AAAA    2001:500:d937::30
m.gtld-servers.net. 172800 IN      AAAA    2001:501:b1f9::30

;; Query time: 40 msec
;; SERVER: 2001:503:ba3e::2:30#53(2001:503:ba3e::2:30)
;; WHEN: Thu Jan 23 13:30:49 EST 2020
;; MSG SIZE  rcvd: 895
```

Figure 2 - .net TLD NS response lacks any A or AAAA RRs

```
; <<> DiG 9.10.6 <<> @f.root-servers.net net ns +nsid
; (2 servers found)
;; global options: +cmd
;; Got answer:
;; ->>HEADER<<- opcode: QUERY, status: NOERROR, id: 58010
;; flags: qr rd; QUERY: 1, ANSWER: 0, AUTHORITY: 13, ADDITIONAL: 1
;; WARNING: recursion requested but not available

;; OPT PSEUDOSECTION:
; EDNS: version: 0, flags:; udp: 1232
; NSID: 49 41 44 2e 63 66 2e 66 2e 72 6f 6f 74 2d 73 65 72 76 65 72 73 2e 6f 72 67 ("IAD.cf.f.root-servers.org")
;; QUESTION SECTION:
;net.                IN      NS

;; AUTHORITY SECTION:
net.                172800  IN      NS      1.gTLD-SeRverS.net.
net.                172800  IN      NS      b.gTLD-SeRverS.net.
net.                172800  IN      NS      c.gTLD-SeRverS.net.
net.                172800  IN      NS      d.gTLD-SeRverS.net.
net.                172800  IN      NS      e.gTLD-SeRverS.net.
net.                172800  IN      NS      f.gTLD-SeRverS.net.
net.                172800  IN      NS      g.gTLD-SeRverS.net.
net.                172800  IN      NS      a.gTLD-SeRverS.net.
net.                172800  IN      NS      h.gTLD-SeRverS.net.
net.                172800  IN      NS      i.gTLD-SeRverS.net.
net.                172800  IN      NS      j.gTLD-SeRverS.net.
net.                172800  IN      NS      k.gTLD-SeRverS.net.
net.                172800  IN      NS      m.gTLD-SeRverS.net.

;; Query time: 7 msec
;; SERVER: 2001:500:2f::f#53(2001:500:2f::f)
;; WHEN: Thu Jan 23 13:31:16 EST 2020
;; MSG SIZE rcvd: 282
```