Passive vs Active Measurement in the DNS

Much of the analysis of DNS behaviours relies on passive measurement in the DNS. The methodology used in this form of analysis has a common pattern:

- select a DNS vantage point (DNS recursive resolver, authoritative name server, root name server)
- capture queries
- make supposition of a DNS behaviour based on the collected data

The issue with this approach is that is tends to assume that all queries have a similar origin, and in general this common origin is assumed to be an end application making a local query via a local stub resolver.

At APNIC we have been interested in an active form of measurement where we launch a large number of DNS queries from a broad sample of end users and observe these queries from a vantage point at an authoritative server.

Our approach uses a unique label in the query where the label has a number of sub-fields that identifies the circumstances of the query and the desired form of response from the authoritative server, casting parts of this query name as a form of micro-code which instructs a particular response from the authoritative server.

> For example, a query generated within this sytem would be:
> 0ds-ud815140d-c13-a1283-s1634709461-icb0a3cd02-0.ap.dotnxdomain.net
>
> In this case the various fields direct the server to add a valid RRSIG record to the response (if the DO bit is set in the query), record the geolocation of the endpoint and the origin AS and the time when the query was passed to the endpoint. In this case the server will respond in either IPv4 or IPv6.
>
> Other fields can be used to select a DNSSEC signing algorithm, the addition of padding elements in the response to generate a larger response, the protocol over which the server will accept queries, and so on. The query can also direct the rcode, including SUCCESS, SERVFAIL and NXDOMAIN.

One of the basic measurements we can conduct within this framework is to calculate the proportion of DNS queries that occur within a short period following the execution of the experiment script on the endpoint ("live queries") and queries whose timestamp indicates some form of query reply is taking place. The long-term average indicates that only some 55% of queries relate to these "live" queries, while the remainder indicate older queries. Over 24-hour periods we see this older query rate take up to 85% of the total query volume for some days. This observation points to a significant level of query stashing and replay in the DNS ecosystem, where the extent of the replay behaviour is now on a par with application-based DNS query volumes.

Our system also allows us to see instances of resolver farm amplification, resolver query loops, resolver shadowing and other forms of pathological resolver behaviour. We have also looked at the issue of the old queries and whether a different Rcode response provokes a different query behaviour.

The generic question posed by this work is: To what extent can analysis of DNS queries gathered from a passive collection of DNS queries inform us of DNS behaviours given that the proportion of various forms of non-live queries is so high?

Geoff Huston,
Joao Damas
APNIC Labs