CT over DNS

Eran Messeri, eranm@google.com

Motivation

Preserving privacy for inclusion proof.

Questions

- Does this achieve the desired goal?
- Will it blend?

Inclusion proof retrieval

Mirror the get-proof-by-hash mechanism from RFC 6962 (Section 4.5).
Query:

TXT <base32 encoded leaf hash>.hash.<domain>
size: 52 characters for the base32-encoded hash, 6 for ‘.hash.’, leaving 195 bytes

The response is (as a TXT RR):

<certificate index in decimal> (size: 1-9 decimal digits).

Inclusion proof retrieval (cont’d)

Then the steps of the Merkle proof are retrieved with:

TXT <level in decimal (0 is the leaf)>.<index in decimal>.<tree_size in decimal>.tree.<domain>

size: ~23 characters assuming a tree size of 10M certs.

Response:

[<node>]+ (size per node: 32 bytes)

Text field will be filled with as many nodes as possible (approx 7)

STH Fetching

Query:

TXT sth.<domain> (size: fixed)

Response:

<tree size in decimal>.<timestamp>.<base 64 root hash>.<base64 signature>
size: ~115 characters for a tree size of 10M, timestamp of 13 digits.

Consistency proof fetching

TXT <entry>.<first>.<second>.sth-consistency.<domain>

Response:

[<node>]+ (size per node: 32 bytes)

Where the first entry is the <entry>’th node in the consistency proof, and the rest of the record is filled with as many entries as will fit (= 255/32 = 7).