

draft-tllq-tsr

Ted Lemon <elemen@apple.com>

Status

- Draft has expired (oops)
- We've learned a bit about the problem, which I presented in San Francisco
- Thinking has gotten a bit clearer

Currently

- TSR is included only in probes, in the authority section, so that it doesn't go into caches
- TSR includes just a time offset
- This code is implemented only in Apple tvOS devices

San Francisco insights

- TSR “conflicts” are not actually conflicts.
- They are stale data.
- A conflict is two hosts claiming authority for same name
- So the model of “conflict resolution” is wrong
- It causes self-inflicted wounds particularly with replication
- The proposal was to include a fingerprint of the SRP update in the TSR message

Subsequent insights

- Fingerprinting the message is the wrong approach.
 - Might have two versions of data from the same host
 - Some data maybe was in different update than other data
- Instead, we can just say who is authoritative for the data
- This doesn't change over time
- We can use the DNS key, or a hash of it, for SRP
- We can use the DNS zone name for actual DNS zones
- When we see different records in an mDNS probe, they aren't seen as in conflict.

More insights

- Need to include TSR for answers, not just authority records
 - This allows us to skip probing if we know we aren't the first to publish (e.g. for SRP replication)
 - This is because if we don't probe, conflicts that would have been resolved by TSR aren't, and show up as actual conflicts when an answer is seen
- Still don't want TSR to show up in caches
- Solution: use an EDNS0 option
- Timing is good
 - We have new stuff to put in the TSR message
 - TSR is only implemented by Apple (right?)
 - Apple devices get updated, so we don't need too much backward compatibility

EDNS0 option format

- Option type, length, obvsly
- Answer section
 - count of records, != ANCOUNT
- Authority section
 - count of records, != NSCOUNT
- Each record is
 - <section index> (what owner name this TSR applies to)
 - TSR
 - <id type> 1 = key (hash?), 2 = zone owner name
- If multiple RRs in a section have same owner name, TSR is only provided for first

Thoughts?

- We are kind of doing research here, but I think we are converging.
- Getting this right will significantly improve behavior when SRP replication and advertising proxy are active together
- (This is the most common use case right now)

- Comments? Thoughts? Bueller?