Workgroup: dnsop
Internet-Draft: draft-tapril-ns2-01
Published: 13 July 2020
Intended Status: Informational
Expires: 14 January 2021
Authors: T. April
Akamai Technologies

Parameterized Nameserver Delegation with NS2 and NS2T

Abstract

Within the DNS, there is no mechanism for authoritative servers to advertise which transport methods they are capable of. If secure transport methods are adopted by authoritative operators, transport signaling would be required to negotiate how authoritative servers would be contacted by resolvers. This document provides two new Resource Record Types, NS2 and NS2T, to facilitate this negotiation by allowing zone owners to signal how the authoritative nameserver(s) for their zone(s) may accept queries.

Discussion Venues

This note is to be removed before publishing as an RFC.

Source for this draft and an issue tracker can be found at https://github.com/timapril/ns2.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 14 January 2021.

Table of Contents

1. Introduction

Resolvers currently rely on the NS records in the parent and child zones to provide and confirm the nameservers that are authoritative for each zone. The Nameserver version 2 (NS2) record extends the functionality of the NS record to include additional information about how authoritative zone information can be queried, whether that be over alternate protocols or by using alternate protocol parameters. The NS2 record may be present at zone cuts but can also redirect resolvers to other nameservers for further redirection via the Nameserver Version 2 Target (NS2T) record, which does not indicate a zone cut.

The NS2 and NS2T records uses the SVCB record format defined in [I-D.draft-ietf-dnsop-svcb-https-00], using a subset of the already defined service parameters as well as new parameters described in this document. Some, but not all, of the existing service parameters will also be available for NS2 and NS2T records. This document will outline the available parameters and their usage.

1.1. Introductory Examples

To introduce the NS2 and NS2T records, this example shows a possible response from an authoritative in the authority section of the DNS response when delegating to another nameserver.

```
example.com.  86400  IN NS2    2 ns2.example.com. ( transports=dot,
                dnsTlsFingerprints=["MIIS987SSLKJ...123===",
                    "MII3SODKSLKJ...456==="] )
example.com.  86400  IN NS2    3 ns3.example.com. ( transports=doh,
                dnsDohURITemplate="https://ns.example.net/q/{?dns}" )
example.com.  86400  IN NS     ns1.example.com.
ns1.example.com.    86400   IN  NS  192.0.2.1
ns2.example.com.    86400   IN  NS  192.0.2.2
ns3.example.com.    86400   IN  NS  192.0.2.3
```

In this example, the authoritative nameserver is delegating to both a DNS-over-TLS and DNS-over-HTTPS service running on ns2.example.net and ns3.example.com respectively, for resolvers that support NS2, and also delegating to ns1.example.com which will serve unencrypted DNS over port 53 for those that do not.

Like in SVCB, NS2 and NS2T also offer the ability to use the Alias form delegation. The example below shows an example where example.com is being delegated with a NS2 AliasForm record which can then be further resolved to locate the authroitative nameserver(s).

```
example.com.  86400  IN NS2 0   ns2.example.net.
example.com.  86400  IN NS     ns1.example.com.
ns1.example.com.    86400   IN  NS  192.0.2.1
```

The example.net authoritative server may return the following NS2T records in response to a query as deirected by the above records.

```
ns2.example.net 3600    IN NS2T ns2.example.org. ( transports=dot,
                dnsTlsFingerprints=["MIIS987SSLKJ...123==="] )
ns2.example.net 3600    IN NS2T ns3.example.org. ( transports=doh,
                ddnsDohURITemplate="https://ns.example.net/q/{?dns}")
```

The avbove records indicate to the client that the authoritative nameservers for zones that Alias to ns2.example.net are ns1.example.org and ns2.example.rg with the configuration provided.

Later sections of this document will go into more detail on the resolution process using these records.

1.2. Goal of the NS2 and NS2T Records

The primary goal of the NS2 and NS2T records is to provide zone owners a way to signal to clients how they may receive queries for the records for which they are authoritative. The NS2 and NS2T records are machine readable, can coexist with NS records in the same zone, and do not break software that does not support them.

NS2 and NS2T are designed to allow child zones to publish NS2 and NS2T records, even without support in the parent zone. Lack of parent zone support for NS2 records may arise for technical or policy reasons.

1.3. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. NS2 and NS2T Record Types

The SVCB record allows for two types of records, the AliasForm and the ServiceForm. The NS2 record takes advantage of both and each will be described below in depth.

2.1. Difference between the records

This document introduces two different resource record types. Both records have the same functionality, with the difference between them being that the NS2 record MUST only be used at a delegation point while the NS2T, is NS2 Target, record does not indicate that the label is being delegated. For example, take the following NS2 record:

```
example.com.  86400  IN NS2 1   ns2.example.net. ( transports=dot )
```

When a client receives the above record, the resolver should send queries for any name under example.com to the nameserver at ns2.example.com unless further delegated. By contrast, when presented with the records below:

```
example.com.  86400  IN NS2 0   customer.example.org.
customer.example.org.  86400  IN NS2T 1   ns2.example.net. ( transports=
```

A resolver trying to resolve a name under example.com would get the first record above from the parent authoritative server, .COM, indicating that the NS2T records found at customer.example.org should be used to locate the authoritative nameservers of example.com. The second record above dictates that the authoritative nameserver from records that have aliased to customer.example.org is ns2.example.net, which only accepts queries over DNS-over-TLS. The primary difference between the two records is that the NS2 record means that anything under the record label should be queried at the delegated server while the NS2T record is just for redirection purposes, and any names under the record's label should still be resolved using the same server unless otherwise delegated.

It should be noted that both NS2 and NS2T records may exist for the same label. Below is an example of this:

```
example.com.  86400  IN NS2 0   c1.example.org.
c1.example.org.  86400  IN NS2T 1   ns2.example.net. ( transports=dot )
c1.example.org.  86400  IN NS2  1   ns3.example.net. ( transports=dot )
test.c1.example.org. 600 IN A 192.0.2.1
```

In the above case, the NS2 record for c1.example.org would only be used when trying to resolve names below c1.example.org. This reason is why when an AliasForm NS2 or NS2T record is encountered, the resolver MUST query for the NS2T record associated with the given name.

Since the NS2 record is indicative of a zone cut while NS2T is not, names MUST NOT have but NS2 and NS2T records at the same time.

2.2. AliasForm Record Type

In order to take full advantage of the AliasForm of NS2 and NS2T, the parent, child and resolver must support these records. When supported, the use of the AliasForm will allow zone owners to delegate their zones to another operator with a single record in the parent. AliasForm NS2 records SHOULD appear in the child zone when used in the parent. If a resolver were to encounter an AliasForm NS2 or NS2T record, it would then resolve the name in the SvcDomainName of the original record using NS2T RR type to receive the either another AliasForm record or a ServiceForm NS2T record.

For example, if the name www.example.com was being resolved, the .com zone may issue a referral by returning the following record:

```
example.com.    86400    IN  NS2     0   customer1.example.net.
```

The above record would indicate to the resolver that in order to obtain the authoritative nameserver records for example.com, the resolver should resolve the RR type NS2T for the name customer1.example.net.

2.2.1. Multiple Service Providers

Some zone owners may wish to use multiple providers to serve their zone, in which case multiple NS2 AliasForm records can be used. In the event that multiple NS2 AliasForm records are encountered, the resolver SHOULD pick one of the records at random. For example, to split traffic between two providers, the zone owner for example.com could have the following NS2 records:

```
example.com.    86400    IN  NS2     0   customer1.example.net.
example.com.    86400    IN  NS2     0   customer1.example.org.
```

DRAFT NOTE: SVCB says that there "SHOULD only have a single RR". This ignores that but keep the randomization part. Section 2.5p1 of SVCB

2.2.2. Loop Prevention

Special care should be taken by both the zone owner and the delegated zone operator to ensure that a lookup loop is not created by having two AliasForm records rely on each other to serve the zone. Doing so may result in a resolution loop, and likely a denial of service. Any clients implementing NS2 and NS2T SHOULD implement a per-resolution limit of how many AliasForm records may be traversed when looking up a delegation to prevent infinite looping. When a loop is detected, like with the handling of CNAME or NS, the server should respond to the client with SERVFAIL.

2.3. ServiceForm Record Type

The ServiceForm of the NS2 and NS2T records are likely to be the more popular of the two. They work the same way as the SVCB or HTTPSSVC records do by providing a priority and list of parameters associated with the SvcDomainName. In addition to being able to identify which protocols are supported by the authoriatative server, the ServiceForm record will also allow providers to operate different protocols on different addresses.

2.3.1. SvcFieldPriority

As defined in the DNS SVCB document [I-D.draft-ietf-dnsop-svcb-httpssvc-00], the SvcFieldPriority values SHOULD be used to dictate the priority when multiple NS2 or NS2T records are returned.

2.3.2. SvcDomainName

As defined in the SVCB document [I-D.draft-ietf-dnsop-svcb-httpssvc-00], the SvcDomainName provides the hostname to which the NS2 or NS2T record refers. The wire format must be the a-label format of the name. When presented, the u-label may be presented, but the the a-label should also be displayed displaying the u-label. The SvcDomainName field MUST be set and MUST NOT be "." for NS2 records. Records with a SvcDomainName of "." SHOULD be discarded.

DRAFT NOTE: Should u-label and a-label be expanded? I'm leaning towards not expanding.

2.3.3. SvcParamKeys

The following DNS SVCB parameters are defined for the NS2 and NS2T ServiceForms.

2.3.3.1. "transports"

The "transports" SvcParamKey defines the list of transports offered by the nameserver named in the SvcDomainName.

The existing "alpn" SvcParamKey was not reused for NS2 and NS2T due to the restriction that the "alpn" SvcParamValues are limited to those defined in the TLS Application-Layer Protocol Negotiation (ALPN) Protocol IDs registry. Plaintext DNS traffic is not, and should not be listed in that registry, but is required to support the transition to encrypted transport in NS2 and NS2T records.

Below is a list of the transport values defined in this document:

*"do53": indicates that a server supports plaintext, unencrypted DNS traffic over UDP or UDP as defined in [RFC1035] and [RFC1034] and the updates to those RFCs.

*"dot": indicates that the server supports encrypted DNS traffic over DNS-over-TLS as defined in [RFC7858].

*"doh": indicates that the server supports encrypted DNS traffic over DNS-over-HTTPS as defined in [RFC8484]. Records that use the DoH service form may be further redirected with HTTPSSVC records in the delegated zone.

*"doq": indicates that the server supports encrypted DNS traffic over DNS over Dedicated QUIC Connections [I-D.draft-huitema-quic-dnsoquic-07]

The order of the keys in the list dictate the order which the nameserver SHOULD be contacted in. The client SHOULD compare the order of available transports with the set of transports it supports to determine how to contact the selected nameserver.

The presentation format of the SvcParamValue is a comma delimited quoted string of the available transport names. The wire format for the SvcParamValue is a string of 16-bit integers representing the TransportKey values as described in the "NS2/NS2T Transport Parameter Registry".

2.3.3.2. "dnsDotEarlyData"

The "dnsDotEarlyData" SvcParamKey indicates if the server will accept requests with TLS 1.3 Early Data as described in [I-D.draft-ghedini-dprive-early-data-01]. If the "dot" transport is enabled on the record but this value is not present, the default value is that the server will not accept early data.

The presentation format of the SvcParamValue is the string "true" or "false". The wire format of the SvcParamValue is to not have the key present or a single octet with the value of 0x00 when early data is not allowed and a 0x01 value when early data is allowed. All other values should be treated as an error and revert the value to the default of not supported.

DRAFT NOTE: Should this be "ns2flags" and just have a 16 bit field for boolean values?

2.3.3.3. "dnsDohURITemplate"

The "dnsDohURITemplate" SvcParamKey defines the URI template to be used for issuing DNS-over-HTTPS queries to the nameserver defined in the record. The host portion of the "dnsDohURITemplate" value SHOULD match the SvcDomainName field.

In the event that the host portion of the "dnsDohURITemplate" SvcParamValues and SvcDomainName field do not match, the SvcDomainName value SHOULD be used for resolving the host and provide host portion of the "dnsDohURITemplate" template SvcParamValue for the TLS ServerNameIndication header and the HTTP Host header. For example, in the below NS2 delegation, the client SHOULD resolve the name ns.example.net and provide the host header and TLS ServerNameIndication header of doh.example.org:

example.com. 86400 IN NS2 3 ns.example.net. (transports=doh, dnsDohURITemplate="https://doh.example.org/q/{?dns}")

The presentation format of the SvcParamValue is a quoted string. The wire form of the SvcParamValue is an octet string of the URI template as defined in [RFC8484].

2.3.3.4. "esniconfig"

The "esnikeys" SvcParamKey is defined in [I-D.draft-ietf-dnsop-svcb-httpssvc-00]. It can be used to provide the ESNI key for DoT, DoH and/or future protocols which may make use of ESNI for session establishment. See [I-D.draft-ietf-dnsop-svcb-httpssvc-00] for the usage, wire, and display formatting for this SvcParamKey.

2.3.3.5. "dnsTlsFingerprints"

The "dnsTlsFingerprints" SvcParamKey is a list of fingerprints for the TLS certificates that may be presented by the nameserver. This record SHOULD match the TLSA record as described in [RFC6698]. Due to bootstrapping concerns, this SvcParamKey has been added to the NS2 record as the TLSA records would only be resolveable after the initial connetion to the delegated nameserver was established. When this field is not present, certificate validation should be performed by either DANE or by traditional TLS certification validation using trusted root certification authorities.

The presentation and wire format of the SvcParamValue is the same as the presentation and wire format described for the TLSA record as defined in [RFC6698], sections 2.1 and 2.2 respectively.

2.4. Deployment Considerations

The NS2 and NS2T records intends to replace the NS record while also adding additional functionality in order to support additional transports for the DNS. Below are discussions of considerations for deployment.

2.4.1. AliasForm and ServiceForm in the Parent

Both the AliasForm and ServiceForm records MUST NOT be returned for the same record when not in delegated zone. In the case where both are present, the ServiceForm MUST be used and the AliasForm ignored.

DRAFT NOTE: This is in direct conflict with SVCB. I'm currently of the opinion that it should stay as it is for reliability reasons. If it is decided that this should contradict SVCB, maybe we should try to change SVCB.

2.4.2. Rollout

When introduced, the NS2 and NS2T record will likely not be supported by the Root or second level operators, and may not be for some time. In the interim, zone owners may place these records into their zones, both for their own zone and any of their child zones. If a resolver supports alternative transports, it MAY, when delegated to another server, issue a query for NS2 or NS2T records and potentially use those records for further query processing.

DRAFT NOTE: Should we include something here that an authoritative MAY include NS2 records in the additionals section of responses to encourage resolvers to upgrade? What about an ECS option from the authority to signal that it is capable of alternat transports?

2.4.3. Availability

If a zone operator removes all NS records before NS2 and NS2T records are implemented by all clients, the availability of their zones will be impacted for the clients that are using non-supporting resolvers. In some cases, this may be a desired quality, but should be noted by zone owners and operators.

2.4.4. Multiple ServiceForm records for the same host or IP address

As described in the "transport" SvcParamKey section above, a host or IP address may support multiple different transport methods. This can be represented in two ways. The first is to list all supported transports in the order of diminishing desire in the same record. The second is to use multiple NS2 or NS2T records.

When those records have different SvcFieldPriority values, as in [I-D.draft-ietf-dnsop-svcb-httpssvc-00], lower-numbered priorities express a higher preference for that record.

NS2 and NS2T records may have multiple values for the "dnsTlsFingerprints" SvcParamKey. Records that are identical other than the "dnsTlsFingerprints" SvcParamValues may be joined together including multiple "dnsTlsFingerprints" as seen in this example:

```
example.com.  86400  IN NS2    2 ns.example.net. ( transports=dot,
                dnsTlsFingerprints="MIIS987SSLKJ...123===" )
example.com.  86400  IN NS2    2 ns.example.net. ( transports=dot,
                dnsTlsFingerprints="MII3SODKSLKJ...456===" )
example.com.  86400  IN NS2    3 ns.example.net. ( transports=doh,
                dnsDohURITemplate="https://dns.example.org/q/{?dns}" )
```

are the same as:

```
example.com.  86400  IN NS2    2 ns.example.net. ( transports=dot,
                dnsTlsFingerprints=["MIIS987SSLKJ...123===",
                    "MII3SODKSLKJ...456==="] )
example.com.  86400  IN NS2    3 ns.example.net. ( transports=doh,
                dnsDohURITemplate="https://dns.example.org/q/{?dns}" )
```

3. Responses with NS2

The NS2 and NS2T records are intended to supersede the NS record. As such, the NS2 records should be included in the response in the following situations, assuming the records exist in the servers zone:

1) If the qtype is for NS2 or NS2T, the server should respond with NS2 or NS2T respectively in the authority section of the response.

2) For queries over unencrypted TCP port 53, any of the NS2 and NS2T records SHOULD be included in the additional section of the response.

3) For queries over unencrypted UDP port 53, any of the NS2 and NS2T records SHOULD be included in the additional section of the response unless doing so would result in a truncated response. For responses that would require truncation, the resolver operator and/or implementor may decide to truncate the response or exclude the records from the response.

4) If encrypted transports are supported on the authority, the any NS2 and NS2T records should be included in the authority section of the respose.

DRAFT NOTE: It is unknown how resolvers will handle multiple authoritative RRTypes in the authority section of the response, leaving 2 and 3 as the additional section until either testing is done or a consensus is determined in DNSOP/DPRIVE.

DRAFT NOTE: Suggesting authority for encrypted transport since it would more closely align with the NS record.

3.1. Response Size Considerations

For latency-conscious zones, the overall packet size of the delegation records from a parent zone to child zone should be taken into account when configuring the NS, NS2 and NS2T records. Resolvers that wish to receive NS2 and NS2T records in response SHOULD advertise and support a buffer size that is as large as possible, ideally 4096 bytes, to allow the authoritative server to respond without truncating whenever possible.

3.2. When to include glue

Like with NS, when a parent is delegating to a child that is in bailiwick, glue records should be included with responses to enable the client to continue to resolve the names.

The ServiceForm version of NS2 returnes sufficent information to the client communicate with the delegated nameserver over a transport other than Do53, but the AliasForm does not. For this reason, the most common use of the AliasForm record would be to alias to a name that is out of bailiwick to the requested zone, which SHOULD be resolveable without relying on the originally queried zone. In the event of an in-bailiwick AliasForm record, the client must either have a pre-agreed upon configuration for the server or attempt opprotunisitic upgrade of the connection to use non-Do53 for the initial setup.

4. DNSSEC and NS2

Like with NS records, the NS2 records in the child zone SHOULD be signed when the zone is DNSSEC signed. The NS2 records that appear in the parent zone are glue and would not be signed, as is the case with NS records.

NS2T records, SHOULD be signed in a zone which is signed by DNSSEC.

5. Privact Considerations

All of the information handled or transmitted by this protocol is public information published in the DNS.

While the records are transmitting public infomration, resolvers which are making use of records may be attemping to keep the information they are querying private from on-path observers. Privacy consious resolvers should query for this record at the apex of a zone when delegated from the parent to help establish an encrypted channel to the authority.

6. Security Considerations

TODO: Fill this section out

6.1. Availability of zones without NS

6.2. Reflection Attacks

6.3. Parsing

6.4. Availability

6.5. Connetion Failures

When a resolver attempts to access nameserver delegated by a NS2 or NST2 record, if a connection error occurs, such as a certificate mismatch or unreachable server, the resolver SHOULD attempt to connect to the other nameservers delegated to until either exhausting the list or the resolver's policy indicates that they should treat the resoltion as failed.

The failure action when failing to resolve a name with NS2/NS2T due to connection erros is dependant on the resolver operators policies. For resolvers which strongly favor privacy, the operators may wish to return a SERVFAIL when the NS2/NS2T resoltion process completes without successfully contacting a delegated nameserver(s) while opportunisitic privacy resolvers may wish to attempt resolution using any NS records that may be present.

7. IANA Considerations

7.1. New registry for NS2 transports

The "NS2/NS2T Transport Parameter Registry" defines the namespace for parameters, including string representations and numeric values. This registry applies to the "transports" DNS SVCB format, currently impacting the NS2 RR Type.

ACTION: create and include a reference to this registry.

7.1.1. Procedure

A registration MUST include the following fields:

*Name: The transport type key name

*TransportKey: A numeric identifier (range 0-65535)

*Meaning: a short description

*Protocol Specification

*Pointer to specification text

7.1.2. Initial Contents

The "NS2/NS2T Transport Parameter Registry" shall initially be populated with the registrations below:

TransportKey	Name	Meaning	Protocol Specification	Reference
0	key0	Reserved	Reserved	(This Document)
1	do53	Unencrypted, Plaintext DNS over UDP or TCP	RFC1035	(This Document)
2	dot	DNS-over-TLS	RFC7858	(This Document)
3	doh	DNS-over-HTTPS	RFC8484	(This Document)
4	doq	DNS over Dedicated QUIC Connections	[DOQ-I-D]	
65280-65534	keyNNNNN	Private Use	Private Use	(This Document)
65535	key65535	Reserved	Reserved	(This Document)

Table 1

7.2. New SvcParamKey Values

This document defines new SvcParamKey values in the "Service Binding (SVCB) Parameter Registry".

SvcParamKey	NAME	Meaning	Reference
TBD1	transports		(This Document)
TBD2	dnsDotEarlyData		(This Document)
TBD3	dnsDohURITemplate		(This Document)
TBD4	dnsTlsFingerprints		(This Document)

Table 2

8. Informative References

[I-D.draft-ietf-dnsop-svcb-https-00]
Schwartz, B., Bishop, M., and E. Nygren, "Service binding and parameter specification via the DNS (DNS SVCB and HTTPS RRs)", Work in Progress, Internet-Draft, draft-ietf-dnsop-svcb-https-00, 12 June 2020, <http://www.ietf.org/internet-drafts/draft-ietf-dnsop-svcb-https-00.txt>.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

[I-D.draft-ietf-dnsop-svcb-httpssvc-00]
Schwartz, B., Bishop, M., and E. Nygren, "Service binding and parameter specification via the DNS (DNS SVCB and HTTPSSVC)", Work in Progress, Internet-Draft, draft-ietf-dnsop-svcb-httpssvc-00, 31 October 2019, <http://www.ietf.org/internet-drafts/draft-ietf-dnsop-svcb-httpssvc-00.txt>.

[RFC1035] Mockapetris, P.V., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[RFC1034] Mockapetris, P.V., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <https://www.rfc-editor.org/info/rfc1034>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[I-D.draft-huitema-quic-dnsoquic-07]
Huitema, C., Shore, M., Mankin, A., Dickinson, S., and J. Iyengar, "Specification of DNS over Dedicated QUIC Connections", Work in Progress, Internet-Draft, draft-

huitema-quic-dnsoquic-07, 7 September 2019, <http://www.ietf.org/internet-drafts/draft-huitema-quic-dnsoquic-07.txt>.

[I-D.draft-ghedini-dprive-early-data-01]
Ghedini, A., "Using Early Data in DNS over TLS", Work in Progress, Internet-Draft, draft-ghedini-dprive-early-data-01, 6 July 2019, <http://www.ietf.org/internet-drafts/draft-ghedini-dprive-early-data-01.txt>.

[RFC6698] Hoffman, P. and J. Schlyter, "The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA", RFC 6698, DOI 10.17487/RFC6698, August 2012, <https://www.rfc-editor.org/info/rfc6698>.

[DOQ-I-D] Huitema, C., Shore, M., Mankin, A., Dickinson, S., and J. Iyengar, "Specification of DNS over Dedicated QUIC Connections", Work in Progress, Internet-Draft, draft-huitema-quic-dnsoquic-07, 7 September 2019, <http://www.ietf.org/internet-drafts/draft-huitema-quic-dnsoquic-07.txt>.

Appendix A. Acknowledgements

Thank you to John Levine, Erik Nygren, Ralf Weber, Jon Reed, Ben Kaduk, Mashooq Muhaimen, Jason Moreau, Jerrod Wiesman, Billy Tiemann, Gordon Marx and Brian Wellington for their comments and suggestions on this draft.

Appendix B. TODO

RFC EDITOR: PLEASE REMOVE THE THIS SECTION PRIOR TO PUBLICATION.

*Discussion about TTLs and what they should be and how that might impact performance

*Discuss the cacheability of the Alias form records.

*Remove all "DRAFT NOTES:" in the document.

*Write a security considerations section

*add prohibition of AliasForm referring to AliasForm

*add out-of-bailiwick requirement for AliasForm

*worked out resolution example including alias form delegation

*DoH URI teamplte does not include post

*If NS2 is authoritative in the parent, does that mean that it will not be a referral anymore? Probably a question for the working group

Appendix C. Change Log

RFC EDITOR: PLEASE REMOVE THE THIS SECTION PRIOR TO PUBLICATION.

pre-00

*Initial draft.

*Wire and Presentation formats for all new SvcParamKeys and SvcParamValues

*IANA considerations first pass

*Added a section about SvcFieldPriority

*Added the "ds" and "dnskey" SvcParamKey options to support the deprecation of the DS in the parent.

*Added notes on DNSSEC signing of NS2

*Removed multi-provider sharding example, with performance measurements the distribution probably wouldn't work

*Reworked the introduction to try and make it easier to parse

*Removed the port fields for each transport option

*Added a discussion about when to include NS2 records.

*Add an example early in the draft. Introduction area

*Add section about port numbers discussion

*collapse udp and tcp to the do53

*added a second record (NS2 and NS2T)

-01

*Removed DS and DNSKEY SvcParamFields. Avoids issues with DNSSEC signing in the parent.

*Removed IPv{4,6}Hints SvcParamFields. There was a discussion on DNSOP about how glue is required

*Updating when to sign the NS2 / NS2T records (removed signing in the parent)

*Attempting to clean up the introduction, goals and motivations of the document

*Adding a privacy considerations section

*Adding more clairty around when to include/expect the NS2/NS2T records

*Adding this note that CNS2 will not be included in this draft

*Prohibiting NS2 and NS2T from existing at the same name

*Making the statement about the parent records being glue and should not be signed

Appendix D. Discussions

Editor Note: Remove this full section before publication.

D.1. Port Numbers

Originally, I had added SvcParamKeys for port numbers for each of the protocols. There was a discussion that resulted in removing the port numbers, since it was added complexity that had little perceived use in the wild. These can be added back if there is a desire to have them. The original author included them as a way to provide the nameserver operator a way to differentiate incoming traffic when using the aliasform with lower TTLs and intelligent responses based on the client IP.

D.2. CNS2

Client NS2, similar to CDS might be a way to provide support for getting NS2 records into the parent zone before going through the registrars, but that might be a tough thing to agree on at this point.

D.3. Second Record Name

Selected NS2T for Nameserver 2 Target since the record defines the target authoritative servers.

~~~
0123456789012345678901234567890123456789012345678901234567890123456789
891
~~~

Author's Address

Tim April
Akamai Technologies

Email: ietf@tapril.net