Authors: C. Huitema   S. Dickinson   A. Mankin
         Private Octopus Inc.   Sinodun IT   Salesforce

DNS over Dedicated QUIC Connections

Abstract

This document describes the use of QUIC to provide transport privacy
for DNS. The encryption provided by QUIC has similar properties to
that provided by TLS, while QUIC transport eliminates the head-of-
line blocking issues inherent with TCP and provides more efficient
packet loss recovery than UDP. DNS over QUIC (DoQ) has privacy
properties similar to DNS over TLS (DoT) specified in RFC7858, and
latency characteristics similar to classic DNS over UDP.

Status of This Memo

Copyright Notice

Section 4.e of the Trust Legal Provisions and are provided without
warranty as described in the Simplified BSD License.

Table of Contents

1. Introduction

Domain Name System (DNS) concepts are specified in "Domain names - concepts and facilities" [RFC1034]. The transmission of DNS queries and responses over UDP and TCP is specified in "Domain names - implementation and specification" [RFC1035].

This document presents a mapping of the DNS protocol over the QUIC transport [RFC9000] [RFC9001]. DNS over QUIC is referred here as DoQ, in line with "DNS Terminology" [I-D.ietf-dnsop-rfc8499bis].

The goals of the DoQ mapping are:

1. Provide the same DNS privacy protection as DNS over TLS (DoT) [RFC7858]. This includes an option for the client to authenticate the server by means of an authentication domain name as specified in "Usage Profiles for DNS over TLS and DNS over DTLS" [RFC8310].

2. Provide an improved level of source address validation for DNS servers compared to classic DNS over UDP.

3. Provide a transport that is not constrained by path MTU limitations on the size of DNS responses it can send.

In order to achieve these goals, and to support ongoing work on encryption of DNS, the scope of this document includes

   *the "stub to recursive resolver" scenario

   *the "recursive resolver to authoritative nameserver" scenario and

   *the "nameserver to nameserver" scenario (mainly used for zone transfers (XFR) [RFC1995], [RFC5936]).

In other words, this document is intended to specify QUIC as a
general purpose transport for DNS.

The specific non-goals of this document are:

1. No attempt is made to evade potential blocking of DNS over QUIC
   traffic by middleboxes.

2. No attempt to support server initiated transactions, which are
   used only in DNS Stateful Operations (DSO) [RFC8490].

Specifying the transmission of an application over QUIC requires
specifying how the application's messages are mapped to QUIC
streams, and generally how the application will use QUIC. This is
done for HTTP in "Hypertext Transfer Protocol Version 3 (HTTP/3)"[I-
D.ietf-quic-http]. The purpose of this document is to define the way
DNS messages can be transmitted over QUIC.

In this document, Section 4 presents the reasoning that guided the
proposed design. Section 5 specifies the actual mapping of DoQ.
Section 6 presents guidelines on the implementation, usage and
deployment of DoQ.

2. Key Words

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT",
"SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this
document are to be interpreted as described in BCP 14 [RFC8174].

3. Document work via GitHub

(RFC EDITOR NOTE: THIS SECTION TO BE REMOVED BEFORE PUBLICATION)The
Github repository for this document is at https://github.com/
huitema/dnsoquic. Proposed text and editorial changes are very much
welcomed there, but any functional changes should always first be
discussed on the IETF DPRIVE WG (dns-privacy) mailing list.

4. Design Considerations

This section and its subsections present the design guidelines that
were used for DoQ. This section is informative in nature.

4.1. Provide DNS Privacy

DoT [RFC7858] defines how to mitigate some of the issues described
in "DNS Privacy Considerations" [RFC9076] by specifying how to
transmit DNS messages over TLS. The "Usage Profiles for DNS over TLS
and DNS over DTLS" [RFC8310] specify Strict and Opportunistic Usage
Profiles for DoT including how stub resolvers can authenticate
recursive resolvers.

QUIC connection setup includes the negotiation of security
parameters using TLS, as specified in "Using TLS to Secure QUIC"
[RFC9001], enabling encryption of the QUIC transport. Transmitting
DNS messages over QUIC will provide essentially the same privacy
protections as DoT [RFC7858] including Strict and Opportunistic
Usage Profiles [RFC8310]. Further discussion on this is provided in
Section 9.

4.2. Design for Minimum Latency

QUIC is specifically designed to reduce protocol-induced delays,
with features such as:

1. Support for 0-RTT data during session resumption.

2. Support for advanced packet loss recovery procedures as
   specified in "QUIC Loss Detection and Congestion Control"
   [RFC9002].

3. Mitigation of head-of-line blocking by allowing parallel
   delivery of data on multiple streams.

This mapping of DNS to QUIC will take advantage of these features in
three ways:

1. Optional support for sending 0-RTT data during session
   resumption (the security and privacy implications of this are
   discussed in later sections).

2. Long-lived QUIC connections over which multiple DNS
   transactions are performed, generating the sustained traffic
   required to benefit from advanced recovery features.

3. Mapping of each DNS Query/Response transaction to a separate
   stream, to mitigate head-of-line blocking. This enables servers
   to respond to queries "out of order". It also enables clients
   to process responses as soon as they arrive, without having to
   wait for in order delivery of responses previously posted by
   the server.

These considerations are reflected in the mapping of DNS traffic to
QUIC streams in Section 5.2.

4.3. No Specific Middlebox Bypass Mechanism

The mapping of DoQ is defined for minimal overhead and maximum
performance. This means a different traffic profile than HTTP3 over
QUIC. This difference can be noted by firewalls and middleboxes.
There may be environments in which HTTP3 over QUIC will be able to
pass through, but DoQ will be blocked by these middle boxes.

4.4. No Server Initiated Transactions

As stated in [Section 1](), this document does not specify support for server initiated transactions within established DoQ connections. That is, only the initiator of the DoQ connection may send queries over the connection.

DSO does support server-initiated transactions within existing connections. However DoQ as defined here does not meet the criteria for an applicable transport for DSO because it does not guarantee in-order delivery of messages, see [Section 4.2]() of [[RFC8490]()].

5. Specifications

5.1. Connection Establishment

DoQ connections are established as described in the QUIC transport specification [[RFC9000]()]. During connection establishment, DoQ support is indicated by selecting the ALPN token "doq" in the crypto handshake.

5.1.1. Draft Version Identification

(RFC EDITOR NOTE: THIS SECTION TO BE REMOVED BEFORE PUBLICATION) Only implementations of the final, published RFC can identify themselves as "doq". Until such an RFC exists, implementations MUST NOT identify themselves using this string.

Implementations of draft versions of the protocol MUST add the string "-" and the corresponding draft number to the identifier. For example, draft-ietf-dprive-dnsoquic-00 is identified using the string "doq-i00".

5.1.2. Port Selection

By default, a DNS server that supports DoQ MUST listen for and accept QUIC connections on the dedicated UDP port TBD (number to be defined in [Section 10]()), unless there is a mutual agreement to use another port.

By default, a DNS client desiring to use DoQ with a particular server MUST establish a QUIC connection to UDP port TBD on the server, unless there is a mutual agreement to use another port.

In order to use a port other than TBD, both clients and servers would need a configuration option in their software.

DoQ connections MUST NOT use UDP port 53. This recommendation against use of port 53 for DoQ is to avoid confusion between DoQ and the use of DNS over UDP [[RFC1035]()].

In the stub to recursive scenario, the use of port 443 as a mutually
agreed alternative port can be operationally beneficial, since port
443 is less likely to be blocked than other ports. Several
mechanisms for stubs to discover recursives offering encrypted
transports, including the use of custom ports, are the subject of
ongoing work.

5.2. Stream Mapping and Usage

The mapping of DNS traffic over QUIC streams takes advantage of the
QUIC stream features detailed in Section 2 of [RFC9000], the QUIC
transport specification.

DNS traffic follows a simple pattern in which the client sends a
query, and the server provides one or more responses (multiple
responses can occur in zone transfers).

The mapping specified here requires that the client selects a
separate QUIC stream for each query. The server then uses the same
stream to provide all the response messages for that query. In order
that multiple responses can be parsed, a 2-octet length field is
used in exactly the same way as the 2-octet length field defined for
DNS over TCP [RFC1035]. The practical result of this is that the
content of each QUIC stream is exactly the same as the content of a
TCP connection that would manage exactly one query.

All DNS messages (queries and responses) sent over DoQ connections
MUST be encoded as a 2-octet length field followed by the message
content as specified in [RFC1035].

The client MUST select the next available client-initiated
bidirectional stream for each subsequent query on a QUIC connection,
in conformance with the QUIC transport specification [RFC9000].

The client MUST send the DNS query over the selected stream, and
MUST indicate through the STREAM FIN mechanism that no further data
will be sent on that stream.

The server MUST send the response(s) on the same stream and MUST
indicate, after the last response, through the STREAM FIN mechanism
that no further data will be sent on that stream.

Therefore, a single client initiated DNS transaction consumes a
single stream. This means that the client's first query occurs on
QUIC stream 0, the second on 4, and so on.

Servers MAY defer processing of a query until the STREAM FIN has
been indicated on the stream selected by the client. Servers and
clients MAY monitor the number of "dangling" streams for which the
expected queries or responses have been received but not the STREAM

FIN. Implementations MAY impose a limit on the number of such
dangling streams. If limits are encountered, implementations MAY
close the connection.

5.2.1. DNS Message IDs

When sending queries over a QUIC connection, the DNS Message ID MUST
be set to zero.

This has implications for proxying DoQ message to and from other
transports. For example, proxies may have to manage the fact that
DoQ can support a larger number of outstanding queries on a single
connection than e.g., DNS over TCP because DoQ is not limited by the
Message ID space.

When forwarding a DNS message from DoQ over another transport, a DNS
Message ID MUST be generated according to the rules of the protocol
that is in use. When forwarding a DNS message from another transport
over DoQ, the Message ID MUST be set to zero.

5.3. DoQ Error Codes

The following error codes are defined for use when abruptly
terminating streams, aborting reading of streams, or immediately
closing connections:

DOQ_NO_ERROR (0x0): No error. This is used when the connection or
   stream needs to be closed, but there is no error to signal.

DOQ_INTERNAL_ERROR (0x1): The DoQ implementation encountered an
   internal error and is incapable of pursuing the transaction or
   the connection.

DOQ_PROTOCOL_ERROR (0x2): The DoQ implementation encountered an
   protocol error and is forcibly aborting the connection.

DOQ_REQUEST_CANCELLED (0x3): A DoQ client uses this to signal that
   it wants to cancel an outstanding transaction.

DOQ_EXCESSIVE_LOAD (0x4): A DoQ implementation uses this to signal
   when closing a connection due to excessive load.

DOQ_ERROR_RESERVED (0xd098ea5e): Alternative error code used for
   tests.

See [Section 10.4](#) for details on registering new error codes.

5.3.1. Transaction Cancellation

In QUIC, sending STOP_SENDING requests that a peer cease transmission on a stream. If a DoQ client wishes to cancel an outstanding request, it MUST issue a QUIC Stop Sending with error code DOQ_REQUEST_CANCELLED. This may be sent at any time but will be ignored if the server has already sent the response. The corresponding DNS transaction MUST be abandoned.

Servers that receive STOP_SENDING act in accordance with [Section 3.5](#) of [[RFC9000](#)]. Servers MAY impose implementation limits on the total number or rate of request cancellations. If limits are encountered, servers MAY close the connection. In this case, servers wanting to help client debugging MAY use the error code DOQ_EXCESSIVE_LOAD. There is always a trade-off between helping good faith clients debug issues and allowing denial-of-service attackers to test server defenses, so depending on circumstances servers might very well chose to send different error codes.

Note that this mechanism provides a way for secondaries to cancel a single zone transfer occurring on a given stream without having to close the QUIC connection.

5.3.2. Transaction Errors

Servers normally complete transactions by sending a DNS response (or responses) on the transaction's stream, including cases where the DNS response indicates a DNS error. For example, a Server Failure (SERVFAIL, [[RFC1035](#)]) SHOULD be notified to the client by sending back a response with the Response Code set to SERVFAIL.

If a server is incapable of sending a DNS response due to an internal error, it SHOULD issue a QUIC Stream Reset. The error code SHOULD be set to DOQ_INTERNAL_ERROR. The corresponding DNS transaction MUST be abandoned. Clients MAY limit the number of unsolicited QUIC Stream Resets received on a connection before choosing to close the connection.

Note that this mechanism provides a way for primaries to abort a single zone transfer occurring on a given stream without having to close the QUIC connection.

5.3.3. Protocol Errors

Other error scenarios can occur due to malformed, incomplete or unexpected messages during a transaction. These include (but are not limited to)

   *a client or server receives a message with a non-zero Message ID

*a client or server receives a STREAM FIN before receiving all the
 bytes for a message indicated in the 2-octet length field

*a client receives a STREAM FIN before receiving all the expected
 responses

*a server receives more than one query on a stream

*a client receives a different number of responses on a stream
 than expected (e.g. multiple responses to a query for an A
 record)

*a client receives a STOP_SENDING request

*the client or server does not indicate the expected STREAM FIN
 after sending requests or responses (see Section 5.2).

*an implementation receives a message containing the edns-tcp-
 keepalive EDNS(0) Option [RFC7828] (see Section 6.5.2)

*a client or a server attempts to open an unidirectional QUIC
 stream

*a server attempts to open a server-initiated bidirectional QUIC
 stream

If a peer encounters such an error condition it is considered a
fatal error. It SHOULD forcibly abort the connection using QUIC's
CONNECTION_CLOSE mechanism, and SHOULD use the DoQ error code
DOQ_PROTOCOL_ERROR.

It is noted that the restrictions on use of the above EDNS(0)
options has implications for proxying message from TCP/DoT/DoH over
DoQ.

5.3.4. Alternative error codes

This specification suggests specific error codes Section 5.3.1,
Section 5.3.2, and Section 5.3.3. These error codes are meant to
facilitates investigation of failures and other incidents. New error
codes may be defined in future versions of DoQ, or registered as
specified in Section 10.4.

Because new error codes can be defined without negotiation, use of
an error code in an unexpected context or receipt of an unknown
error code MUST be treated as equivalent to DOQ_NO_ERROR.

Implementations MAY wish to test the support for the error code
extension mechanism by using error codes not listed in this
document, or they MAY use DOQ_ERROR_RESERVED.

5.4. Connection Management

[Section 10](#) of [[RFC9000](#)], the QUIC transport specification, specifies that connections can be closed in three ways:

  *idle timeout

  *immediate close

  *stateless reset

Clients and servers implementing DoQ SHOULD negotiate use of the idle timeout. Closing on idle timeout is done without any packet exchange, which minimizes protocol overhead. Per [Section 10.1](#) of [[RFC9000](#)], the QUIC transport specification, the effective value of the idle timeout is computed as the minimum of the values advertised by the two endpoints. Practical considerations on setting the idle timeout are discussed in [Section 6.5.2](#).

Clients SHOULD monitor the idle time incurred on their connection to the server, defined by the time spent since the last packet from the server has been received. When a client prepares to send a new DNS query to the server, it will check whether the idle time is sufficient lower than the idle timer. If it is, the client will send the DNS query over the existing connection. If not, the client will establish a new connection and send the query over that connection.

Clients MAY discard their connections to the server before the idle timeout expires. A client that has outstanding queries SHOULD close the connection explicitly using QUIC's CONNECTION_CLOSE mechanism and the DoQ error code DOQ_NO_ERROR.

Clients and servers MAY close the connection for a variety of other reasons, indicated using QUIC's CONNECTION_CLOSE. Client and servers that send packets over a connection discarded by their peer MAY receive a stateless reset indication. If a connection fails, all the in progress transaction on that connection MUST be abandoned.

5.5. Session Resumption and 0-RTT

A client MAY take advantage of the session resumption mechanisms supported by QUIC transport [[RFC9000](#)] and QUIC TLS [[RFC9001](#)]. Clients SHOULD consider potential privacy issues associated with session resumption before deciding to use this mechanism. These privacy issues are detailed in [Section 9.2](#) and [Section 9.1](#), and the implementation considerations are discussed in [Section 6.5.3](#).

The 0-RTT mechanism SHOULD NOT be used to send DNS requests that are not "replayable" transactions. In this specification, only transactions that have an OPCODE of QUERY or NOTIFY are considered

replayable and MAY be sent in 0-RTT data. See Appendix A for a
detailed discussion of why NOTIFY is included here.

Servers MUST NOT execute non replayable transactions received in 0-
RTT data. Servers MUST adopt one of the following behaviors:

  *Queue the offending transaction and only execute it after the
   QUIC handshake has been completed, as defined in Section 4.1.1 of
   [RFC9001].

  *Reply to the offending transaction with a response code REFUSED
   and an Extended DNS Error Code (EDE) "Too Early", see Section
   10.3.

  *Close the connection with the error code DOQ_PROTOCOL_ERROR.

5.6. Message Sizes

DoQ Queries and Responses are sent on QUIC streams, which in theory
can carry up to 2^{62} bytes. However, DNS messages are restricted in
practice to a maximum size of 65535 bytes. This maximum size is
enforced by the use of a two-octet message length field in DNS over
TCP [RFC1035] and DNS over TLS [RFC7858], and by the definition of
the "application/dns-message" for DNS over HTTP [RFC8484]. DoQ
enforces the same restriction.

The Extension Mechanisms for DNS (EDNS) [RFC6891] allow peers to
specify the UDP message size. This parameter is ignored by DoQ. DoQ
implementations always assume that the maximum message size is 65535
bytes.

6. Implementation Requirements

6.1. Authentication

For the stub to recursive resolver scenario, the authentication
requirements are the same as described in DoT [RFC7858] and "Usage
Profiles for DNS over TLS and DNS over DTLS" [RFC8310]. [RFC8932]
states that DNS privacy services SHOULD provide credentials that
clients can use to authenticate the server. Given this, and to align
with the authentication model for DoH, DoQ stubs SHOULD use a Strict
authentication profile. Client authentication for the encrypted stub
to recursive scenario is not described in any DNS RFC.

For zone transfer, the requirements are the same as described in
[RFC9103].

For the recursive resolver to authoritative nameserver scenario,
authentication requirements are unspecified at the time of writing
and are the subject on ongoing work in the DPRIVE WG.

6.2. Fallback to Other Protocols on Connection Failure

If the establishment of the DoQ connection fails, clients MAY
attempt to fall back to DoT and then potentially clear text, as
specified in DoT [RFC7858] and "Usage Profiles for DNS over TLS and
DNS over DTLS" [RFC8310], depending on their privacy profile.

DNS clients SHOULD remember server IP addresses that don't support
DoQ. Timeouts, connection refusals, and QUIC handshake failures are
valid indicators that a server does not support DoQ. Clients SHOULD
NOT attempt DoQ queries to a server that does not support DoQ for a
reasonable period (such as one hour per server). DNS clients
following an out-of-band key-pinned privacy profile ([RFC7858]) MAY
be more aggressive about retrying DoQ connection failures.

6.3. Address Validation

Section 8 of [RFC9000], the QUIC transport specification, defines
Address Validation procedures to avoid servers being used in address
amplification attacks. DoQ implementations MUST conform to this
specification, which limits the worst case amplification to a factor
3.

DoQ implementations SHOULD consider configuring servers to use the
Address Validation using Retry Packets procedure defined in
Section 8.1.2 of [RFC9000], the QUIC transport specification. This
procedure imposes a 1-RTT delay for verifying the return routability
of the source address of a client, similar to the DNS Cookies
mechanism [RFC7873].

DoQ implementations that configure Address Validation using Retry
Packets SHOULD implement the Address Validation for Future
Connections procedure defined in Section 8.1.3 of [RFC9000], the
QUIC transport specification. This defines how servers can send
NEW_TOKEN frames to clients after the client address is validated,
in order to avoid the 1-RTT penalty during subsequent connections by
the client from the same address.

6.4. Padding

Implementations SHOULD protect against the traffic analysis attacks
described in Section 9.5 by the judicious injection of padding. This
could be done either by padding individual DNS messages using the
EDNS(0) Padding Option [RFC7830] and by padding QUIC packets (see
Section 8.6 of [RFC9000], the QUIC transport specification.

In theory, padding at the QUIC level could result in better
performance for the equivalent protection, because the amount of
padding can take into account non-DNS frames such as
acknowledgeemnts or flow control updates, and also because QUIC

packets can carry multiple DNS messages. However, applications can only control the amount of padding in QUIC packets if the implementation of QUIC exposes adequate APIs. This leads to the following recommendation:

  *if the implementation of QUIC exposes APIs to set a padding policy, DNS over QUIC SHOULD use that API to align the packet length to a small set of fixed sizes, aligned with the recommendations of the "Padding Policies for Extension Mechanisms for DNS (EDNS(0))" [RFC8467].

  *if padding at the QUIC level is not available or not used, DNS over QUIC MUST ensure that all DNS queries and responses are padded to a small set of fixed sizes, using the EDNS padding extension as specified in "Padding Policies for Extension Mechanisms for DNS (EDNS(0))" [RFC8467].

6.5. Connection Handling

"DNS Transport over TCP - Implementation Requirements" [RFC7766] provides updated guidance on DNS over TCP, some of which is applicable to DoQ. This section provides similar advice on connection handling for DoQ.

6.5.1. Connection Reuse

Historic implementations of DNS clients are known to open and close TCP connections for each DNS query. To amortise connection setup costs, both clients and servers SHOULD support connection reuse by sending multiple queries and responses over a single persistent QUIC connection.

In order to achieve performance on par with UDP, DNS clients SHOULD send their queries concurrently over the QUIC streams on a QUIC connection. That is, when a DNS client sends multiple queries to a server over a QUIC connection, it SHOULD NOT wait for an outstanding reply before sending the next query.

6.5.2. Resource Management

Proper management of established and idle connections is important to the healthy operation of a DNS server.

An implementation of DoQ SHOULD follow best practices similar to those specified for DNS over TCP [RFC7766], in particular with regard to:

  *Concurrent Connections (Section 6.2.2 of [RFC7766], updated by Section 6.4 of [RFC9103])

*Security Considerations ([Section 10](#) of [[RFC7766](#)])

Failure to do so may lead to resource exhaustion and denial of service.

Clients that want to maintain long duration DoQ connections SHOULD use the idle timeout mechanisms defined in [Section 10.1](#) of [[RFC9000](#)], the QUIC transport specification. Clients and servers MUST NOT send the edns-tcp-keepalive EDNS(0) Option [[RFC7828](#)] in any messages sent on a DoQ connection (because it is specific to the use of TCP/TLS as a transport).

This document does not make specific recommendations for timeout values on idle connections. Clients and servers should reuse and/or close connections depending on the level of available resources. Timeouts may be longer during periods of low activity and shorter during periods of high activity.

6.5.3. Using 0-RTT and Session Resumption

Using 0-RTT for DNS over QUIC has many compelling advantages. Clients can establish connections and send queries without incurring a connection delay. Servers can thus negotiate low values of the connection timers, which reduces the total number of connections that they need to manage. They can do that because the clients that use 0-RTT will not incur latency penalties if new connections are required for a query.

Session resumption and 0-RTT data transmission create privacy risks detailed in detailed in [Section 9.2](#) and [Section 9.1](#). The following recommendations are meant to reduce the privacy risks while enjoying the performance benefits of 0-RTT data, with the restriction specified in [Section 5.5](#).

Clients SHOULD use resumption tickets only once, as specified in Appendix C.4 to [[RFC8446](#)]. By default, clients SHOULD NOT use session resumption if the client's connectivity has changed.

Clients could receive address validation tokens from the server using the NEW_TOKEN mechanism; see [Section 8](#) of [[RFC9000](#)]. The associated tracking risks are mentioned in [Section 9.3](#). Clients SHOULD only use the address validation tokens when they are also using session resumption, thus avoiding additional tracking risks.

Servers SHOULD issue session resumption tickets with a sufficiently long life time (e.g., 6 hours), so that clients are not tempted to either keep connection alive or frequently poll the server to renew session resumption tickets. Servers SHOULD implement the anti-replay mechanisms specified in [Section 8](#) of [[RFC8446](#)].

6.5.4. Controlling Connection Migration For Privacy

DoQ implementation might consider using the connection migration features defined in Section 9 of [RFC9000]. These features enable connections to continue operating as the client's connectivity changes. As detailed in Section 9.4, these features trade off privacy for latency. By default, clients SHOULD be configured to prioritise privacy and start new sessions if their connectivity changes.

6.6. Processing Queries in Parallel

As specified in Section 7 of [RFC7766] "DNS Transport over TCP - Implementation Requirements", resolvers are RECOMMENDED to support the preparing of responses in parallel and sending them out of order. In DoQ, they do that by sending responses on their specific stream as soon as possible, without waiting for availability of responses for previously opened streams.

6.7. Zone transfer

[RFC9103] specifies zone transfer over TLS (XoT) and includes updates to [RFC1995] (IXFR), [RFC5936] (AXFR) and [RFC7766]. Considerations relating to the re-use of XoT connections described there apply analogously to zone transfers performed using DoQ connections. For example:

*DoQ servers MUST be able to handle multiple concurrent IXFR requests on a single QUIC connection

*DoQ servers MUST be able to handle multiple concurrent AXFR requests on a single QUIC connection

*DoQ implementations SHOULD

-use the same QUIC connection for both AXFR and IXFR requests to the same primary

-pipeline such requests (if they pipeline XFR requests in general) and MAY intermingle them

-send the response(s) for each request as soon as they are available i.e. responses MAY be sent intermingled

6.8. Flow Control Mechanisms

Servers and Clients manage flow control using the mechanisms defined in Section 4 of [RFC9000]. These mechanisms allow clients and servers to specify how many streams can be created, how much data can be sent on a stream, and how much data can be sent on the union

of all streams. For DNS over QUIC, controlling how many streams are created allows servers to control how many new requests the client can send on a given connection.

Flow control exists to protect endpoint resources. For servers, global and per-stream flow control limits control how much data can be sent by clients. The same mechanisms allow clients to control how much data can be sent by servers. Values that are too small will unnecessarily limit performance. Values that are too large might expose endpoints to overload or memory exhaustion. Implementations or deployments will need to adjust flow control limits to balance these concerns. In particular, zone transfer implementations will need to control these limits carefully to ensure both large and concurrent zone transfers are well managed.

Initial values of parameters control how many requests and how much data can be sent by clients and servers at the beginning of the connection. These values are specified in transport parameters exchanged during the connection handshake. The parameter values received in the initial connection also control how many requests and how much data can be sent by clients using 0-RTT data in a resumed connection. Using too small values of these initial parameters would restrict the usefulness of allowing 0-RTT data.

7. Implementation Status

(RFC EDITOR NOTE: THIS SECTION TO BE REMOVED BEFORE PUBLICATION) This section records the status of known implementations of the protocol defined by this specification at the time of posting of this Internet-Draft, and is based on a proposal described in [RFC7942].

1. AdGuard launched a DoQ recursive resolver service in December 2020. They have released a suite of open source tools that support DoQ:

   1. AdGuard C++ DNS libraries A DNS proxy library that supports all existing DNS protocols including DNS-over-TLS, DNS-over-HTTPS, DNSCrypt and DNS-over-QUIC (experimental).

   2. DNS Proxy A simple DNS proxy server that supports all existing DNS protocols including DNS-over-TLS, DNS-over-HTTPS, DNSCrypt, and DNS-over-QUIC. Moreover, it can work as a DNS-over-HTTPS, DNS-over-TLS or DNS-over-QUIC server.

   3. CoreDNS fork for AdGuard DNS Includes DNS-over-QUIC server-side support.

4. dnslookup Simple command line utility to make DNS lookups. Supports all known DNS protocols: plain DNS, DoH, DoT, DoQ, DNSCrypt.

2. Quicdoq Quicdoq is a simple open source implementation of DoQ. It is written in C, based on Picoquic.

3. Flamethrower is an open source DNS performance and functional testing utility written in C++ that has an experimental implementation of DoQ.

4. aioquic is an implementation of QUIC in Python. It includes example client and server for DoQ.

7.1. Performance Measurements

To our knowledge, no benchmarking studies comparing DoT, DoH and DoQ are published yet. However anecdotal evidence from the AdGuard DoQ recursive resolver deployment indicates that it performs well compared to the other encrypted protocols, particularly in mobile environments. Reasons given for this include that DoQ

*Uses less bandwidth due to a more efficient handshake (and due to less per message overhead when compared to DoH).

*Performs better in mobile environments due to the increased resilience to packet loss

*Can maintain connections as users move between mobile networks via its connection management

8. Security Considerations

The security considerations of DoQ should be comparable to those of DoT [RFC7858].

9. Privacy Considerations

The general considerations of encrypted transports provided in "DNS Privacy Considerations" [RFC9076] apply to DoQ. The specific considerations provided there do not differ between DoT and DoQ, and are not discussed further here. Similarly, "Recommendations for DNS Privacy Service Operators" [RFC8932] (which covers operational, policy, and security considerations for DNS privacy services) is also applicable to DoQ services.

QUIC incorporates the mechanisms of TLS 1.3 [RFC8446] and this enables QUIC transmission of "0-RTT" data. This can provide interesting latency gains, but it raises two concerns:

1. Adversaries could replay the 0-RTT data and infer its content from the behavior of the receiving server.

2. The 0-RTT mechanism relies on TLS session resumption, which can provide linkability between successive client sessions.

These issues are developed in Section 9.1 and Section 9.2.

9.1. Privacy Issues With 0-RTT data

The 0-RTT data can be replayed by adversaries. That data may trigger queries by a recursive resolver to authoritative resolvers. Adversaries may be able to pick a time at which the recursive resolver outgoing traffic is observable, and thus find out what name was queried for in the 0-RTT data.

This risk is in fact a subset of the general problem of observing the behavior of the recursive resolver discussed in "DNS Privacy Considerations" [RFC9076]. The attack is partially mitigated by reducing the observability of this traffic. The mandatory replay protection mechanisms in TLS 1.3 [RFC8446] limit but do not eliminate the risk of replay. 0-RTT packets can only be replayed within a narrow window, which is only wide enough to account for variations in clock skew and network transmission.

The recommendation for TLS 1.3 [RFC8446] is that the capability to use 0-RTT data should be turned off by default, and only enabled if the user clearly understands the associated risks. In our case, allowing 0-RTT data provides significant performance gains, and we are concerned that a recommendation to not use it would simply be ignored. Instead, we provide a set of practical recommendations in Section 5.5 and Section 6.5.3.

The prevention on allowing replayable transactions in 0-RTT data expressed in Section 5.5 blocks the most obvious risks of replay attacks, as it only allows for transactions that will not change the long term state of the server.

The attacks described above apply to the stub resolver to recursive resolver scenario, but similar attacks might be envisaged in the recursive resolver to authoritative resolver scenario, and the same mitigations apply.

9.2. Privacy Issues With Session Resumption

The QUIC session resumption mechanism reduces the cost of re-establishing sessions and enables 0-RTT data. There is a linkability issue associated with session resumption, if the same resumption token is used several times. Attackers on path between client and server could observe repeated usage of the token and use that to track the client over time or over multiple locations.

The session resumption mechanism allows servers to correlate the resumed sessions with the initial sessions, and thus to track the client. This creates a virtual long duration session. The series of queries in that session can be used by the server to identify the client. Servers can most probably do that already if the client address remains constant, but session resumption tickets also enable tracking after changes of the client's address.

The recommendations in Section 6.5.3 are designed to mitigate these risks. Using session tickets only once mitigates the risk of tracking by third parties. Refusing to resume a session if addresses change mitigates the risk of tracking by the server.

The privacy trade-offs here may be context specific. Stub resolvers will have a strong motivation to prefer privacy over latency since they often change location. However, recursive resolvers that use a small set of static IP addresses are more likely to prefer the reduced latency provided by session resumption and may consider this a valid reason to use resumption tickets even if the IP address changed between sessions.

Encrypted zone transfer (RFC9103) explicitly does not attempt to hide the identity of the parties involved in the transfer, but at the same time such transfers are not particularly latency sensitive. This means that applications supporting zone transfers may decide to apply the same protections as stub to recursive applications.

9.3. Privacy Issues With Address Validation Tokens

QUIC specifies address validation mechanisms in Section 8 of [RFC9000]. Use of an address validation token allows QUIC servers to avoid an extra RTT for new connections. Address validation tokens are typically tied to an IP address. QUIC clients normally only use these tokens when setting a new connection from a previously used address. However, due to the prevalence of NAT, clients are not always aware that they are using a new address. There is a linkability risk if clients mistakenly use address validation tokens after unknowingly moving to a new location.

The recommendations in Section 6.5.3 mitigates this risk by tying the usage of the NEW_TOKEN to that of session resumption.

9.4. Privacy Issues With Long Duration Sessions

A potential alternative to session resumption is the use of long
duration sessions: if a session remains open for a long time, new
queries can be sent without incurring connection establishment
delays. It is worth pointing out that the two solutions have similar
privacy characteristics. Session resumption may allow servers to
keep track of the IP addresses of clients, but long duration
sessions have the same effect.

In particular, a DoQ implementation might take advantage of the
connection migration features of QUIC to maintain a session even if
the client's connectivity changes, for example if the client
migrates from a Wi-Fi connection to a cellular network connection,
and then to another Wi-Fi connection. The server would be able to
track the client location by monitoring the succession of IP
addresses used by the long duration connection.

The recommendation in Section 6.5.4 mitigates the privacy concerns
related to long duration sessions using multiple client addresses.

9.5. Traffic Analysis

Even though QUIC packets are encrypted, adversaries can gain
information from observing packet lengths, in both queries and
responses, as well as packet timing. Many DNS requests are emitted
by web browsers. Loading a specific web page may require resolving
dozen of DNS names. If an application adopts a simple mapping of one
query or response per packet, or "one QUIC STREAM frame per packet",
then the succession of packet lengths may provide enough information
to identify the requested site.

Implementations SHOULD use the mechanisms defined in Section 6.4 to
mitigate this attack.

10. IANA Considerations

10.1. Registration of DoQ Identification String

This document creates a new registration for the identification of
DoQ in the "Application Layer Protocol Negotiation (ALPN) Protocol
IDs" registry [RFC7301].

The "doq" string identifies DoQ:

Protocol: DoQ

Identification Sequence: 0x64 0x6F 0x71 ("doq")

Specification: This document

10.2. Reservation of Dedicated Port

Port 853 is currently reserved for 'DNS query-response protocol run over TLS/DTLS' [RFC7858]. However, the specification for DNS over DTLS (DoD) [RFC8094] is experimental, limited to stub to resolver, and no implementations or deployments currently exist to our knowledge (even though several years have passed since the specification was published).

This specification proposes to additionally reserve the use of port 853 for DoQ. QUIC was designed to be able to co-exist with other protocols on the same port, including DTLS , see Section 17.2 of [RFC9000].

IANA is requested to add the following value to the "Service Name and Transport Protocol Port Number Registry" in the System Range. The registry for that range requires IETF Review or IESG Approval [RFC6335].

Service Name: dns-over-quic

Port Number: 853

Transport Protocol(s): UDP

Assignee: IESG

Contact: IETF Chair

Description: DNS query-response protocol run over QUIC

Reference: This document

10.2.1. Port number 784 for experimentations

(RFC EDITOR NOTE: THIS SECTION TO BE REMOVED BEFORE PUBLICATION) Early experiments MAY use port 784. This port is marked in the IANA registry as unassigned.

(Note that version in -02 of this draft experiments were directed to use port 8853.)

10.3. Reservation of Extended DNS Error Code Too Early

IANA is requested to add the following value to the Extended DNS Error Codes registry [RFC8914]:

INFO-CODE: TBD

Purpose: Too Early

Reference:

This document

10.4. DNS over QUIC Error Codes Registry

IANA [SHALL add/has added] a registry for "DNS over QUIC Error
Codes" on the "Domain Name System (DNS) Parameters" web page.

The "DNS over QUIC Error Codes" registry governs a 62-bit space.
This space is split into three regions that are governed by
different policies:

  *Permanent registrations for values between 0x00 and 0x3f (in
   hexadecimal; inclusive), which are assigned using Standards
   Action or IESG Approval as defined in Section 4.9 and Section
   4.10 of [RFC8126]

  *Permanent registrations for values larger than 0x3f, which are
   assigned using the Specification Required policy ([RFC8126])

  *Provisonal registrations for values larger than 0x3f, which
   require Expert Review, as defined in Section 4.5 of [RFC8126].

Provisional reservations share the range of values larger than 0x3f
with some permanent registrations. This is by design, to enable
conversion of provisional registrations into permanent registrations
without requiring changes in deployed systems. (This design is
aligned with the principles set in Section 22 of [RFC9000].)

Registrations in this registry MUST include the following fields:

Value: The assigned codepoint.

Status: "Permanent" or "Provisional".

Contact: Contact details for the registrant.

Notes: Supplementary notes about the registration.

In addition, permanent registrations MUST include:

Error: A short mnemonic for the parameter.

Specification: A reference to a publicly available specification
   for the value (optional for provisional registrations).

Description: A brief description of the error code semantics, which
   MAY be a summary if a specification reference is provided.

Provisional registrations of codepoints are intended to allow for private use and experimentation with extensions to DNS over QUIC. However, provisional registrations could be reclaimed and reassigned for another purpose. In addition to the parameters listed above, provisional registrations MUST include:

Date: The date of last update to the registration.

A request to update the date on any provisional registration can be made without review from the designated expert(s).

The initial contents of this registry are shown in Table 1.

Value	Error	Description	Specification
0x0	DOQ_NO_ERROR	No error	Section 5.3
0x1	DOQ_INTERNAL_ERROR	Implementation error	Section 5.3
0x2	DOQ_PROTOCOL_ERROR	Generic protocol violation	Section 5.3
0x3	DOQ_REQUEST_CANCELLED	Request cancelled by client	Section 5.3
0x4	DOQ_EXCESSIVE_LOAD	Closing a connection for excessive load	Section 5.3
0xd098ea5e	DOQ_ERROR_RESERVED	Alternative error code used for tests	Section 5.3

Table 1: Initial DNS over QUIC Error Codes Entries

11. Acknowledgements

This document liberally borrows text from the HTTP-3 specification [I-D.ietf-quic-http] edited by Mike Bishop, and from the DoT specification [RFC7858] authored by Zi Hu, Liang Zhu, John Heidemann, Allison Mankin, Duane Wessels, and Paul Hoffman.

The privacy issue with 0-RTT data and session resumption were analyzed by Daniel Kahn Gillmor (DKG) in a message to the IETF "DPRIVE" working group [DNS0RTT].

Thanks to Tony Finch for an extensive review of the initial version of this draft, and to Robert Evans for the discussion of 0-RTT privacy issues. Reviews by Paul Hoffman and Martin Thomson and interoperability tests conducted by Stephane Bortzmeyer helped improve the definition of the protocol.

12. References

12.1. Normative References

[I-D.ietf-dnsop-rfc8499bis]

Hoffman, P. and K. Fujiwara, "DNS Terminology", Work in Progress, Internet-Draft, draft-ietf-dnsop-rfc8499bis-03, 28 September 2021, <https://www.ietf.org/archive/id/draft-ietf-dnsop-rfc8499bis-03.txt>.

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

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

[RFC1995] Ohta, M., "Incremental Zone Transfer in DNS", RFC 1995, DOI 10.17487/RFC1995, August 1996, <https://www.rfc-editor.org/info/rfc1995>.

[RFC5936] Lewis, E. and A. Hoenes, Ed., "DNS Zone Transfer Protocol (AXFR)", RFC 5936, DOI 10.17487/RFC5936, June 2010, <https://www.rfc-editor.org/info/rfc5936>.

[RFC6891] Damas, J., Graff, M., and P. Vixie, "Extension Mechanisms for DNS (EDNS(0))", STD 75, RFC 6891, DOI 10.17487/RFC6891, April 2013, <https://www.rfc-editor.org/info/rfc6891>.

[RFC7301] Friedl, S., Popov, A., Langley, A., and E. Stephan, "Transport Layer Security (TLS) Application-Layer Protocol Negotiation Extension", RFC 7301, DOI 10.17487/RFC7301, July 2014, <https://www.rfc-editor.org/info/rfc7301>.

[RFC7766] Dickinson, J., Dickinson, S., Bellis, R., Mankin, A., and D. Wessels, "DNS Transport over TCP - Implementation Requirements", RFC 7766, DOI 10.17487/RFC7766, March 2016, <https://www.rfc-editor.org/info/rfc7766>.

[RFC7828] Wouters, P., Abley, J., Dickinson, S., and R. Bellis, "The edns-tcp-keepalive EDNS0 Option", RFC 7828, DOI 10.17487/RFC7828, April 2016, <https://www.rfc-editor.org/info/rfc7828>.

[RFC7830] Mayrhofer, A., "The EDNS(0) Padding Option", RFC 7830, DOI 10.17487/RFC7830, May 2016, <https://www.rfc-editor.org/info/rfc7830>.

[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>.

[RFC7873] Eastlake 3rd, D. and M. Andrews, "Domain Name System
(DNS) Cookies", RFC 7873, DOI 10.17487/RFC7873, May 2016,
<https://www.rfc-editor.org/info/rfc7873>.

[RFC8094] Reddy, T., Wing, D., and P. Patil, "DNS over Datagram
Transport Layer Security (DTLS)", RFC 8094, DOI 10.17487/
RFC8094, February 2017, <https://www.rfc-editor.org/info/
rfc8094>.

[RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for
Writing an IANA Considerations Section in RFCs", BCP 26,
RFC 8126, DOI 10.17487/RFC8126, June 2017, <https://
www.rfc-editor.org/info/rfc8126>.

[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>.

[RFC8310] Dickinson, S., Gillmor, D., and T. Reddy, "Usage Profiles
for DNS over TLS and DNS over DTLS", RFC 8310, DOI
10.17487/RFC8310, March 2018, <https://www.rfc-
editor.org/info/rfc8310>.

[RFC8467] Mayrhofer, A., "Padding Policies for Extension Mechanisms
for DNS (EDNS(0))", RFC 8467, DOI 10.17487/RFC8467,
October 2018, <https://www.rfc-editor.org/info/rfc8467>.

[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>.

[RFC8914] Kumari, W., Hunt, E., Arends, R., Hardaker, W., and D.
Lawrence, "Extended DNS Errors", RFC 8914, DOI 10.17487/
RFC8914, October 2020, <https://www.rfc-editor.org/info/
rfc8914>.

[RFC9000] Iyengar, J., Ed. and M. Thomson, Ed., "QUIC: A UDP-Based
Multiplexed and Secure Transport", RFC 9000, DOI
10.17487/RFC9000, May 2021, <https://www.rfc-editor.org/
info/rfc9000>.

[RFC9001] Thomson, M., Ed. and S. Turner, Ed., "Using TLS to Secure
QUIC", RFC 9001, DOI 10.17487/RFC9001, May 2021,
<https://www.rfc-editor.org/info/rfc9001>.

[RFC9103] Toorop, W., Dickinson, S., Sahib, S., Aras, P., and A.
Mankin, "DNS Zone Transfer over TLS", RFC 9103, DOI

                 10.17487/RFC9103, August 2021, <https://www.rfc-
                 editor.org/info/rfc9103>.

12.2. Informative References

   [DNS0RTT]   Kahn Gillmor, D., "DNS + 0-RTT", Message to DNS-Privacy
                WG mailing list, 6 April 2016, <https://www.ietf.org/
                mail-archive/web/dns-privacy/current/msg01276.html>.

   [I-D.ietf-quic-http] Bishop, M., "Hypertext Transfer Protocol
                Version 3 (HTTP/3)", Work in Progress, Internet-Draft,
                draft-ietf-quic-http-34, 2 February 2021, <https://
                www.ietf.org/archive/id/draft-ietf-quic-http-34.txt>.

   [RFC1996]   Vixie, P., "A Mechanism for Prompt Notification of Zone
                Changes (DNS NOTIFY)", RFC 1996, DOI 10.17487/RFC1996,
                August 1996, <https://www.rfc-editor.org/info/rfc1996>.

   [RFC6335]   Cotton, M., Eggert, L., Touch, J., Westerlund, M., and S.
                Cheshire, "Internet Assigned Numbers Authority (IANA)
                Procedures for the Management of the Service Name and
                Transport Protocol Port Number Registry", BCP 165, RFC
                6335, DOI 10.17487/RFC6335, August 2011, <https://
                www.rfc-editor.org/info/rfc6335>.

   [RFC7942]   Sheffer, Y. and A. Farrel, "Improving Awareness of
                Running Code: The Implementation Status Section", BCP
                205, RFC 7942, DOI 10.17487/RFC7942, July 2016, <https://
                www.rfc-editor.org/info/rfc7942>.

   [RFC8446]   Rescorla, E., "The Transport Layer Security (TLS)
                Protocol Version 1.3", RFC 8446, DOI 10.17487/RFC8446,
                August 2018, <https://www.rfc-editor.org/info/rfc8446>.

   [RFC8490]   Bellis, R., Cheshire, S., Dickinson, J., Dickinson, S.,
                Lemon, T., and T. Pusateri, "DNS Stateful Operations",
                RFC 8490, DOI 10.17487/RFC8490, March 2019, <https://
                www.rfc-editor.org/info/rfc8490>.

   [RFC8932]   Dickinson, S., Overeinder, B., van Rijswijk-Deij, R.,
                and A. Mankin, "Recommendations for DNS Privacy Service

Operators", BCP 232, RFC 8932, DOI 10.17487/RFC8932, October 2020, <https://www.rfc-editor.org/info/rfc8932>.

[RFC9002] Iyengar, J., Ed. and I. Swett, Ed., "QUIC Loss Detection and Congestion Control", RFC 9002, DOI 10.17487/RFC9002, May 2021, <https://www.rfc-editor.org/info/rfc9002>.

[RFC9076] Wicinski, T., Ed., "DNS Privacy Considerations", RFC 9076, DOI 10.17487/RFC9076, July 2021, <https://www.rfc-editor.org/info/rfc9076>.

Appendix A. The NOTIFY Service

This appendix discusses why it is considered acceptable to send NOTIFY (see [RFC1996]) in 0-RTT data.

Section 5.5 says "The 0-RTT mechanism SHOULD NOT be used to send DNS requests that are not "replayable" transactions". This specification supports sending a NOTIFY in 0-RTT data because although a NOTIFY technically changes the state of the receiving server, the effect of replaying NOTIFYs has negligible impact in practice.

NOTIFY messages prompt a secondary to either send an SOA query or an XFR request to the primary on the basis that a newer version of the zone is available. It has long been recognized that NOTIFYs can be forged and, in theory, used to cause a secondary to send repeated unnecessary requests to the primary. For this reason, most implementations have some form of throttling of the SOA/XFR queries triggered by the receipt of one or more NOTIFYs.

[RFC9103] describes the privacy risks associated with both NOTIFY and SOA queries and does not include addressing those risks within the scope of encrypting zone transfers. Given this, the privacy benefit of using DoQ for NOTIFY is not clear - but for the same reason, sending NOTIFY as 0-RTT data has no privacy risk above that of sending it using cleartext DNS.

Appendix B. Notable Changes From Previous Versions

(RFC EDITOR NOTE: THIS SECTION TO BE REMOVED BEFORE PUBLICATION)

B.1. Stream Mapping Incompatibility With Draft-02

Versions prior to -02 of this specification proposed a simpler mapping scheme of queries and responses to QUIc stream, which omitted the 2 byte length field and supported only a single response on a given stream. The more complex mapping in Section 5.2 was adopted to specifically cater for XFR support, however it breaks compatibility with earlier versions.

Authors' Addresses

Christian Huitema
Private Octopus Inc.
427 Golfcourse Rd
Friday Harbor, WA 98250
United States of America

Email: huitema@huitema.net

Sara Dickinson
Sinodun IT
Oxford Science Park
Oxford
OX4 4GA
United Kingdom

Email: sara@sinodun.com

Allison Mankin
Salesforce

Email: allison.mankin@gmail.com