

Workgroup: QUIC
Internet-Draft:
draft-ietf-quic-load-balancers-13
Published: 28 March 2022
Intended Status: Standards Track
Expires: 29 September 2022
Authors: M. Duke N. Banks C. Huitema
 Google Microsoft Private Octopus Inc.
QUIC-LB: Generating Routable QUIC Connection IDs

Abstract

QUIC address migration allows clients to change their IP address while maintaining connection state. To reduce the ability of an observer to link two IP addresses, clients and servers use new connection IDs when they communicate via different client addresses. This poses a problem for traditional "layer-4" load balancers that route packets via the IP address and port 4-tuple. This specification provides a standardized means of securely encoding routing information in the server's connection IDs so that a properly configured load balancer can route packets with migrated addresses correctly. As it proposes a structured connection ID format, it also provides a means of connection IDs self-encoding their length to aid some hardware offloads.

Note to Readers

Discussion of this document takes place on the QUIC Working Group mailing list (quic@ietf.org), which is archived at <https://mailarchive.ietf.org/arch/browse/quic/>.

Source for this draft and an issue tracker can be found at <https://github.com/quicwg/load-balancers>.

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 29 September 2022.

Copyright Notice

Copyright (c) 2022 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. [Introduction](#)
 - 1.1. [Terminology](#)
 - 1.2. [Notation](#)
2. [First CID octet](#)
 - 2.1. [Config Rotation](#)
 - 2.2. [Configuration Failover](#)
 - 2.3. [Length Self-Description](#)
 - 2.4. [Format](#)
3. [Load Balancing Preliminaries](#)
 - 3.1. [Unroutable Connection IDs](#)
 - 3.2. [Fallback Algorithms](#)
 - 3.3. [Server ID Allocation](#)
4. [Server ID Encoding in Connection IDs](#)
 - 4.1. [CID format](#)
 - 4.2. [Configuration Agent Actions](#)
 - 4.3. [Server Actions](#)
 - 4.3.1. [Special Case: Single Pass Encryption](#)
 - 4.3.2. [General Case: Four-Pass Encryption](#)
 - 4.4. [Load Balancer Actions](#)
 - 4.4.1. [Special Case: Single Pass Encryption](#)
 - 4.4.2. [General Case: Four-Pass Encryption](#)
5. [Per-connection state](#)
6. [Additional Use Cases](#)
 - 6.1. [Load balancer chains](#)
 - 6.2. [Moving connections between servers](#)
7. [Version Invariance of QUIC-LB](#)
8. [Security Considerations](#)
 - 8.1. [Attackers not between the load balancer and server](#)
 - 8.2. [Attackers between the load balancer and server](#)
 - 8.3. [Multiple Configuration IDs](#)

- [8.4. Limited configuration scope](#)
- [8.5. Stateless Reset Oracle](#)
- [8.6. Connection ID Entropy](#)
- [9. IANA Considerations](#)
- [10. References](#)
 - [10.1. Normative References](#)
 - [10.2. Informative References](#)
- [Appendix A. QUIC-LB YANG Model](#)
 - [A.1. Tree Diagram](#)
- [Appendix B. Load Balancer Test Vectors](#)
 - [B.1. Unencrypted CIDs](#)
 - [B.2. Encrypted CIDs](#)
- [Appendix C. Interoperability with DTLS over UDP](#)
 - [C.1. DTLS 1.0 and 1.2](#)
 - [C.2. DTLS 1.3](#)
 - [C.3. Future Versions of DTLS](#)
- [Appendix D. Acknowledgments](#)
- [Appendix E. Change Log](#)
 - [E.1. since draft-ietf-quic-load-balancers-12](#)
 - [E.2. since draft-ietf-quic-load-balancers-11](#)
 - [E.3. since draft-ietf-quic-load-balancers-10](#)
 - [E.4. since draft-ietf-quic-load-balancers-09](#)
 - [E.5. since draft-ietf-quic-load-balancers-08](#)
 - [E.6. since draft-ietf-quic-load-balancers-07](#)
 - [E.7. since draft-ietf-quic-load-balancers-06](#)
 - [E.8. since draft-ietf-quic-load-balancers-05](#)
 - [E.9. since draft-ietf-quic-load-balancers-04](#)
 - [E.10. since draft-ietf-quic-load-balancers-03](#)
 - [E.11. since draft-ietf-quic-load-balancers-02](#)
 - [E.12. since draft-ietf-quic-load-balancers-01](#)
 - [E.13. since draft-ietf-quic-load-balancers-00](#)
 - [E.14. Since draft-duke-quic-load-balancers-06](#)
 - [E.15. Since draft-duke-quic-load-balancers-05](#)
 - [E.16. Since draft-duke-quic-load-balancers-04](#)
 - [E.17. Since draft-duke-quic-load-balancers-03](#)
 - [E.18. Since draft-duke-quic-load-balancers-02](#)
 - [E.19. Since draft-duke-quic-load-balancers-01](#)
 - [E.20. Since draft-duke-quic-load-balancers-00](#)
- [Authors' Addresses](#)

1. Introduction

QUIC packets [[RFC9000](#)] usually contain a connection ID to allow endpoints to associate packets with different address/port 4-tuples to the same connection context. This feature makes connections robust in the event of NAT rebinding. QUIC endpoints usually designate the connection ID which peers use to address packets. Server-generated connection IDs create a potential need for out-of-band communication to support QUIC.

QUIC allows servers (or load balancers) to designate an initial connection ID to encode useful routing information for load balancers. It also encourages servers, in packets protected by cryptography, to provide additional connection IDs to the client. This allows clients that know they are going to change IP address or port to use a separate connection ID on the new path, thus reducing linkability as clients move through the world.

There is a tension between the requirements to provide routing information and mitigate linkability. Ultimately, because new connection IDs are in protected packets, they must be generated at the server if the load balancer does not have access to the connection keys. However, it is the load balancer that has the context necessary to generate a connection ID that encodes useful routing information. In the absence of any shared state between load balancer and server, the load balancer must maintain a relatively expensive table of server-generated connection IDs, and will not route packets correctly if they use a connection ID that was originally communicated in a protected `NEW_CONNECTION_ID` frame.

This specification provides common algorithms for encoding the server mapping in a connection ID given some shared parameters. The mapping is generally only discoverable by observers that have the parameters, preserving unlinkability as much as possible.

As this document proposes a structured QUIC Connection ID, it also proposes a system for self-encoding connection ID length in all packets, so that crypto offload can efficiently obtain key information.

While this document describes a small set of configuration parameters to make the server mapping intelligible, the means of distributing these parameters between load balancers, servers, and other trusted intermediaries is out of its scope. There are numerous well-known infrastructures for distribution of configuration.

1.1. Terminology

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 RFC 2119 [[RFC2119](#)].

In this document, these words will appear with that interpretation only when in ALL CAPS. Lower case uses of these words are not to be interpreted as carrying significance described in RFC 2119.

In this document, "client" and "server" refer to the endpoints of a QUIC connection unless otherwise indicated. A "load balancer" is an intermediary for that connection that does not possess QUIC connection keys, but it may rewrite IP addresses or conduct other IP

or UDP processing. A "configuration agent" is the entity that determines the QUIC-LB configuration parameters for the network and leverages some system to distribute that configuration.

Note that stateful load balancers that act as proxies, by terminating a QUIC connection with the client and then retrieving data from the server using QUIC or another protocol, are treated as a server with respect to this specification.

For brevity, "Connection ID" will often be abbreviated as "CID".

1.2. Notation

All wire formats will be depicted using the notation defined in Section 1.3 of [RFC9000]. There is one addition: the function `len()` refers to the length of a field which can serve as a limit on a different field, so that the lengths of two fields can be concisely defined as limited to a sum, for example:

`x(A..B) y(C..B-len(x))`

indicates that `x` can be of any length between `A` and `B`, and `y` can be of any length between `C` and `B` provided that `(len(x) + len(y))` does not exceed `B`.

The example below illustrates the basic framework:

```
Example Structure {
  One-bit Field (1),
  7-bit Field with Fixed Value (7) = 61,
  Field with Variable-Length Integer (i),
  Arbitrary-Length Field (..),
  Variable-Length Field (8..24),
  Variable-Length Field with Dynamic Limit (8..24-len(Variable-Length Fi
  Field With Minimum Length (16..),
  Field With Maximum Length (..128),
  [Optional Field (64)],
  Repeated Field (8) ...,
}
```

Figure 1: Example Format

2. First CID octet

The first octet of a Connection ID is reserved for two special purposes, one mandatory (config rotation) and one optional (length self-description).

Subsequent sections of this document refer to the contents of this octet as the "first octet."

2.1. Config Rotation

The first two bits of any connection ID MUST encode an identifier for the configuration that the connection ID uses. This enables incremental deployment of new QUIC-LB settings (e.g., keys).

When new configuration is distributed to servers, there will be a transition period when connection IDs reflecting old and new configuration coexist in the network. The rotation bits allow load balancers to apply the correct routing algorithm and parameters to incoming packets.

Configuration Agents SHOULD deliver new configurations to load balancers before doing so to servers, so that load balancers are ready to process CIDs using the new parameters when they arrive.

A Configuration Agent SHOULD NOT use a codepoint to represent a new configuration until it takes precautions to make sure that all connections using CIDs with an old configuration at that codepoint have closed or transitioned.

Servers MUST NOT generate new connection IDs using an old configuration after receiving a new one from the configuration agent. Servers MUST send NEW_CONNECTION_ID frames that provide CIDs using the new configuration, and retire CIDs using the old configuration using the "Retire Prior To" field of that frame.

It also possible to use these bits for more long-lived distinction of different configurations, but this has privacy implications (see [Section 8.3](#)).

2.2. Configuration Failover

If a server has not received a valid QUIC-LB configuration, and believes that low-state, Connection-ID aware load balancers are in the path, it SHOULD generate connection IDs with the config rotation bits set to '11' and SHOULD use the "disable_active_migration" transport parameter in all new QUIC connections. It SHOULD NOT send NEW_CONNECTION_ID frames with new values.

A load balancer that sees a connection ID with config rotation bits set to '11' MUST revert to 5-tuple routing. These connection IDs may be of any length; however, see [Section 8.6](#) for limits on this length.

2.3. Length Self-Description

Local hardware cryptographic offload devices may accelerate QUIC servers by receiving keys from the QUIC implementation indexed to the connection ID. However, on physical devices operating multiple QUIC servers, it is impractical to efficiently lookup these keys if the connection ID does not self-encode its own length.

Note that this is a function of particular server devices and is irrelevant to load balancers. As such, load balancers MAY omit this from their configuration. However, the remaining 6 bits in the first octet of the Connection ID are reserved to express the length of the following connection ID, not including the first octet.

A server not using this functionality SHOULD make the six bits appear to be random.

2.4. Format

```
First Octet {  
  Config Rotation (2),  
  CID Len or Random Bits (6),  
}
```

Figure 2: First Octet Format

The first octet has the following fields:

Config Rotation: Indicates the configuration used to interpret the CID.

CID Len or Random Bits: Length Self-Description (if applicable), or random bits otherwise. Encodes the length of the Connection ID following the First Octet.

3. Load Balancing Preliminaries

In QUIC-LB, load balancers do not generate individual connection IDs for servers. Instead, they communicate the parameters of an algorithm to generate routable connection IDs.

The algorithms differ in the complexity of configuration at both load balancer and server. Increasing complexity improves obfuscation of the server mapping.

This section describes three participants: the configuration agent, the load balancer, and the server. For any given QUIC-LB configuration that enables connection-ID-aware load balancing, there

must be a choice of (1) routing algorithm, (2) server ID allocation strategy, and (3) algorithm parameters.

Fundamentally, servers generate connection IDs that encode their server ID. Load balancers decode the server ID from the CID in incoming packets to route to the correct server.

There are situations where a server pool might be operating two or more routing algorithms or parameter sets simultaneously. The load balancer uses the first two bits of the connection ID to multiplex incoming DCIDs over these schemes (see [Section 2.1](#)).

3.1. Unroutable Connection IDs

QUIC-LB servers will generate Connection IDs that are decodable to extract a server ID in accordance with a specified algorithm and parameters. However, QUIC often uses client-generated Connection IDs prior to receiving a packet from the server.

These client-generated CIDs might not conform to the expectations of the routing algorithm and therefore not be routable by the load balancer. Those that are not routable are "unroutable DCIDs" and receive similar treatment regardless of why they're unroutable:

- *The config rotation bits ([Section 2.1](#)) may not correspond to an active configuration. Note: a packet with a DCID that indicates 5-tuple routing (see [Section 2.2](#)) is always routable.

- *The DCID might not be long enough for the decoder to process.

- *The extracted server mapping might not correspond to an active server.

All other DCIDs are routable.

Load balancers MUST forward packets with routable DCIDs to a server in accordance with the chosen routing algorithm. Exception: if the load balancer can parse the QUIC packet and makes a routing decision depending on the contents (e.g., the SNI in a TLS client hello), it MAY route in accordance with this instead. However, load balancers MUST always route long header packets it cannot parse in accordance with the DCID (see [Section 7](#)).

Load balancers SHOULD drop short header packets with unroutable DCIDs.

When forwarding a packet with a long header and unroutable DCID, load balancers MUST use a fallback algorithm as specified in [Section 3.2](#).

Load balancers MAY drop packets with long headers and unroutable DCIDs if and only if it knows that the encoded QUIC version does not allow an unroutable DCID in a packet with that signature. For example, a load balancer can safely drop a QUIC version 1 Handshake packet with an unroutable DCID, as a version 1 Handshake packet sent to a QUIC-LB routable server will always have a server-generated routable CID. The prohibition against dropping packets with long headers remains for unknown QUIC versions.

Furthermore, while the load balancer function MUST NOT drop packets, the device might implement other security policies, outside the scope of this specification, that might force a drop.

Servers that receive packets with unroutable CIDs MUST use the available mechanisms to induce the client to use a routable CID in future packets. In QUIC version 1, this requires using a routable CID in the Source CID field of server-generated long headers.

3.2. Fallback Algorithms

There are conditions described below where a load balancer routes a packet using a "fallback algorithm." It can choose any algorithm, without coordination with the servers, but the algorithm SHOULD be deterministic over short time scales so that related packets go to the same server. The design of this algorithm SHOULD consider the version-invariant properties of QUIC described in [[RFC8999](#)] to maximize its robustness to future versions of QUIC.

A fallback algorithm MUST NOT make the routing behavior dependent on any bits in the first octet of the QUIC packet header, except the first bit, which indicates a long header. All other bits are QUIC version-dependent and intermediaries SHOULD NOT base their design on version-specific templates.

For example, one fallback algorithm might convert a unroutable DCID to an integer and divided by the number of servers, with the modulus used to forward the packet. The number of servers is usually consistent on the time scale of a QUIC connection handshake. Another might simply hash the address/port 4-tuple. See also [Section 7](#).

3.3. Server ID Allocation

Load Balancer configurations include a mapping of server IDs to forwarding addresses. The corresponding server configurations contain one or more unique server IDs.

The configuration agent chooses a server ID length for each configuration that MUST be at least one octet.

A QUIC-LB configuration MAY significantly over-provision the server ID space (i.e., provide far more codepoints than there are servers) to increase the probability that a randomly generated Destination Connection ID is unroutable.

The configuration agent SHOULD provide a means for servers to express the number of server IDs it can usefully employ, because a single routing address actually corresponds to multiple server entities (see [Section 6.1](#)).

Conceptually, each configuration has its own set of server ID allocations, though two static configurations with identical server ID lengths MAY use a common allocation between them.

A server encodes one of its assigned server IDs in any CID it generates using the relevant configuration.

4. Server ID Encoding in Connection IDs

4.1. CID format

All connection IDs use the following format:

```
QUIC-LB Connection ID {  
    First Octet (8),  
    Server ID (8..152-len(Nonce)),  
    Nonce (32..152-len(Server ID)),  
}
```

Figure 3: CID Format

4.2. Configuration Agent Actions

The configuration agent assigns a server ID to every server in its pool in accordance with [Section 3.3](#), and determines a server ID length (in octets) sufficiently large to encode all server IDs, including potential future servers.

Each configuration specifies the length of the Server ID and Nonce fields, with limits defined for each algorithm.

Optionally, it also defines a 16-octet key. Note that failure to define a key means that observers can determine the assigned server of any connection, significantly increasing the linkability of QUIC address migration.

The nonce length MUST be at least 4 octets. The server ID length MUST be at least 1 octet.

As QUIC version 1 limits connection IDs to 20 octets, the server ID and nonce lengths MUST sum to 19 octets or less.

4.3. Server Actions

The server writes the first octet and its server ID into their respective fields.

If there is no key in the configuration, the server MUST fill the Nonce field with bytes that appear to be random. If there is a key, the server fills the nonce field with a nonce of its choosing. See [Section 8.6](#) for details.

The server MAY append additional bytes to the connection ID, up to the limit specified in that version of QUIC, for its own use. These bytes MUST NOT provide observers with any information that could link two connection IDs to the same connection, client, or server. In particular, all servers using a configuration MUST consistently add the same length to each connection ID, to preserve the linkability objectives of QUIC-LB. Any additional bytes SHOULD appear random unless individual servers are not distinguishable (e.g. any server using that configuration appends identical bytes to every connection ID).

If there is no key in the configuration, the Connection ID is complete. Otherwise, there are further steps, as described in the two following subsections.

Encryption below uses the AES-128-ECB cipher. Future standards could add new algorithms that use other ciphers to provide cryptographic agility in accordance with [[RFC7696](#)]. QUIC-LB implementations SHOULD be extensible to support new algorithms.

4.3.1. Special Case: Single Pass Encryption

When the nonce length and server ID length sum to exactly 16 octets, the server MUST use a single-pass encryption algorithm. All connection ID octets except the first form an AES-ECB block. This block is encrypted once, and the result forms the second through seventeenth most significant bytes of the connection ID.

4.3.2. General Case: Four-Pass Encryption

Any other field length requires four passes for encryption and at least three for decryption. To understand this algorithm, it is useful to define four functions that minimize the amount of bit-shifting necessary in the event that there are an odd number of octets.

The `expand_left()` function outputs 16 octets, with its first argument in the most significant bits, its second argument in the least significant byte, and zeros in all other positions. Thus,

```
expand_left(0xaaba3c, 0x13) = 0xaaba3c00000000000000000000000013
```

`expand_right()` is similar, except that the second argument is in the most significant byte, and the first argument is in the least significant bits. Therefore,

```
expand_right(0xaaba3c, 0x13) = 0x1300000000000000000000000000aaba3c
```

Similarly, `truncate_left()` and `truncate_right()` take the most significant and least significant bits, respectively, from a ciphertext. For example, to take 28 bits of a ciphertext:

```
truncate_left(0x2094842ca49256198c2deaa0ba53caa0, 28) = 0x2094842  
truncate_right(0x2094842ca49256198c2deaa0ba53caa0, 28) = 0xa53caa0
```

The example at the end of this section helps to clarify the steps described below.

1. The server concatenates the server ID and nonce to create `plaintext_CID`.
2. The server splits `plaintext_CID` into components `left_0` and `right_0` of equal length, splitting an odd octet in half if necessary. For example, `0x7040b81b55ccf3` would split into a `left_0` of `0x7040b81` and `right_0` of `0xb55ccf3`.
3. Encrypt the result of `expand_left(left_0)` to obtain a ciphertext.
4. XOR the least significant bits of the ciphertext with `right_0` to form `right_1`.

Thus steps 3 and 4 can be expressed as `right_1 = right_0 ^ truncate_right(AES_ECB(key, expand_left(left_0, 0x01)), len(right_0))`

5. Repeat steps 3 and 4, but use them to compute `left_1` by expanding and encrypting `right_1` with the most significant octet as `0x02` and XOR the results with `left_0`.

```
left_1 = left_0 ^ truncate_left( AES_ECB(key,  
expand_right(right_1, 0x02)), len(left_0))
```

6. Repeat steps 3 and 4, but use them to compute right_2 by expanding and encrypting left_1 with the least significant octet as 0x03 and XOR the results with right_1.

```
right_2 = right_1 ^ truncate_right( AES_ECB(key,  
expand_left(left_1, 0x03)), len(right_1))
```

7. Repeat steps 3 and 4, but use them to compute left_2 by expanding and encrypting right_2 with the most significant octet as 0x04 and XOR the results with left_1.

```
left_2 = left_1 ^ truncate_left( AES_ECB(key,  
expand_right(right_2, 0x04)), len(left_1))
```

8. The server concatenates left_2 with right_2 to form the ciphertext CID, which it appends to the first octet.

The following example executes the steps for the provided inputs. Note that the plaintext is of odd octet length, so the middle octet will be split evenly left_0 and right_0.

the AES-ECB key and extracts the server ID from the most significant bytes of the resulting plaintext.

4.4.2. General Case: Four-Pass Encryption

First, split the ciphertext CID (excluding the first octet) into its equal-length components `left_2` and `right_2`. Then follow the process below:

```
left_1 = left_2 ^ truncate_left(AES_ECB(key, expand_right(right_2), 0x04))
right_1 = right_2 ^ truncate_right(AES_ECB(key, expand_left(left_1, 0x03))
left_0 = left_1 ^ truncate_left(AES_ECB(key, expand_right(right_1), 0x02))
```

As the load balancer has no need for the nonce, it can conclude after 3 passes as long as the server ID is entirely contained in `left_0` (i.e., the nonce is at least as large as the server ID). If the server ID is longer, a fourth pass is necessary:

```
right_0 = right_1 ^ truncate_right(AES_ECB(key, expand_left(left_0,
0x01)))
```

and the load balancer has to concatenate `left_0` and `right_0` to obtain the complete server ID.

5. Per-connection state

QUIC-LB requires no per-connection state at the load balancer. The load balancer can extract the server ID from the connection ID of each incoming packet and route that packet accordingly.

However, once the routing decision has been made, the load balancer MAY associate the 4-tuple with the decision. This has two advantages:

- *The load balancer only extracts the server ID once per incoming 4-tuple. When the CID is encrypted, this substantially reduces computational load.

- *Incoming Stateless Reset packets and ICMP messages are easily routed to the correct origin server.

In addition to the increased state requirements, however, load balancers cannot detect the `CONNECTION_CLOSE` frame to indicate the end of the connection, so they rely on a timeout to delete connection state. There are numerous considerations around setting such a timeout.

In the event a connection ends, freeing an IP and port, and a different connection migrates to that IP and port before the timeout,

the load balancer will misroute the different connection's packets to the original server. A short timeout limits the likelihood of such a misrouting.

Furthermore, if a short timeout causes premature deletion of state, the routing is easily recoverable by decoding an incoming Connection ID. However, a short timeout also reduces the chance that an incoming Stateless Reset is correctly routed.

Servers MAY implement the technique described in Section 14.4.1 of [\[RFC9000\]](#) in case the load balancer is stateless, to increase the likelihood a Source Connection ID is included in ICMP responses to Path Maximum Transmission Unit (PMTU) probes. Load balancers MAY parse the echoed packet to extract the Source Connection ID, if it contains a QUIC long header, and extract the Server ID as if it were in a Destination CID.

6. Additional Use Cases

This section discusses considerations for some deployment scenarios not implied by the specification above.

6.1. Load balancer chains

Some network architectures may have multiple tiers of low-state load balancers, where a first tier of devices makes a routing decision to the next tier, and so on, until packets reach the server. Although QUIC-LB is not explicitly designed for this use case, it is possible to support it.

If each load balancer is assigned a range of server IDs that is a subset of the range of IDs assigned to devices that are closer to the client, then the first devices to process an incoming packet can extract the server ID and then map it to the correct forwarding address. Note that this solution is extensible to arbitrarily large numbers of load-balancing tiers, as the maximum server ID space is quite large.

If the number of necessary server IDs per next hop is uniform, a simple implementation would use successively longer server IDs at each tier of load balancing, and the server configuration would match the last tier. The forward load balancers would simply treat the least significant bits of the server ID as part of the nonce.

6.2. Moving connections between servers

Some deployments may transparently move a connection from one server to another. The means of transferring connection state between servers is out of scope of this document.

To support a handover, a server involved in the transition could issue CIDs that map to the new server via a `NEW_CONNECTION_ID` frame, and retire CIDs associated with the new server using the "Retire Prior To" field in that frame.

Alternately, if the old server is going offline, the load balancer could simply map its server ID to the new server's address.

7. Version Invariance of QUIC-LB

The server ID encodings, and requirements for their handling, are designed to be QUIC version independent (see [\[RFC8999\]](#)). A QUIC-LB load balancer will generally not require changes as servers deploy new versions of QUIC. However, there are several unlikely future design decisions that could impact the operation of QUIC-LB.

The maximum Connection ID length could be below the minimum necessary for one or more encoding algorithms.

[Section 3.1](#) provides guidance about how load balancers should handle unroutable DCIDs. This guidance, and the implementation of an algorithm to handle these DCIDs, rests on some assumptions:

- *Incoming short headers do not contain DCIDs that are client-generated.
- *The use of client-generated incoming DCIDs does not persist beyond a few round trips in the connection.
- *While the client is using DCIDs it generated, some exposed fields (IP address, UDP port, client-generated destination Connection ID) remain constant for all packets sent on the same connection.

While this document does not update the commitments in [\[RFC8999\]](#), the additional assumptions are minimal and narrowly scoped, and provide a likely set of constants that load balancers can use with minimal risk of version- dependence.

If these assumptions are invalid, this specification is likely to lead to loss of packets that contain unroutable DCIDs, and in extreme cases connection failure.

Some load balancers might inspect elements of the Server Name Indication (SNI) extension in the TLS Client Hello to make a routing decision. Note that the format and cryptographic protection of this information may change in future versions or extensions of TLS or QUIC, and therefore this functionality is inherently not version-invariant. See also [Section 3.1](#) for other considerations about this case. Note that an SNI-aware load balancer, faced with an unknown QUIC version, might misdirect initial packets to the wrong tenant.

While inefficient, this preserves the ability for tenants to deploy new versions provided they have an out-of-band means of providing a connection ID for the client to use.

8. Security Considerations

QUIC-LB is intended to prevent linkability. Attacks would therefore attempt to subvert this purpose.

Note that without a key for the encoding, QUIC-LB makes no attempt to obscure the server mapping, and therefore does not address these concerns. Without a key, QUIC-LB merely allows consistent CID encoding for compatibility across a network infrastructure, which makes QUIC robust to NAT rebinding. Servers that are encoding their server ID without a key algorithm SHOULD only use it to generate new CIDs for the Server Initial Packet and SHOULD NOT send CIDs in QUIC NEW_CONNECTION_ID frames, except that it sends one new Connection ID in the event of config rotation [Section 2.1](#). Doing so might falsely suggest to the client that said CIDs were generated in a secure fashion.

A linkability attack would find some means of determining that two connection IDs route to the same server. As described above, there is no scheme that strictly prevents linkability for all traffic patterns, and therefore efforts to frustrate any analysis of server ID encoding have diminishing returns.

8.1. Attackers not between the load balancer and server

Any attacker might open a connection to the server infrastructure and aggressively simulate migration to obtain a large sample of IDs that map to the same server. It could then apply analytical techniques to try to obtain the server encoding.

An encrypted encoding provides robust protection against this. An unencrypted one provides none.

Were this analysis to obtain the server encoding, then on-path observers might apply this analysis to correlating different client IP addresses.

8.2. Attackers between the load balancer and server

Attackers in this privileged position are intrinsically able to map two connection IDs to the same server. The QUIC-LB algorithms do prevent the linkage of two connection IDs to the same individual connection if servers make reasonable selections when generating new IDs for that connection.

8.3. Multiple Configuration IDs

During the period in which there are multiple deployed configuration IDs (see [Section 2.1](#)), there is a slight increase in linkability. The server space is effectively divided into segments with CIDs that have different config rotation bits. Entities that manage servers SHOULD strive to minimize these periods by quickly deploying new configurations across the server pool.

8.4. Limited configuration scope

A simple deployment of QUIC-LB in a cloud provider might use the same global QUIC-LB configuration across all its load balancers that route to customer servers. An attacker could then simply become a customer, obtain the configuration, and then extract server IDs of other customers' connections at will.

To avoid this, the configuration agent SHOULD issue QUIC-LB configurations to mutually distrustful servers that have different keys for encryption algorithms. In many cases, the load balancers can distinguish these configurations by external IP address.

However, assigning multiple entities to an IP address is complimentary with concealing DNS requests (e.g., DoH [[RFC8484](#)]) and the TLS Server Name Indicator (SNI) ([\[I-D.ietf-tls-esni\]](#)) to obscure the ultimate destination of traffic. While the load balancer's fallback algorithm ([Section 3.2](#)) can use the SNI to make a routing decision on the first packet, there are three ways to route subsequent packets:

- *all co-tenants can use the same QUIC-LB configuration, leaking the server mapping to each other as described above;
- *co-tenants can be issued one of up to three configurations distinguished by the config rotation bits ([Section 2.1](#)), exposing information about the target domain to the entire network; or
- *tenants can use 4-tuple routing in their CIDs (in which case they SHOULD disable migration in their connections), which neutralizes the value of QUIC-LB but preserves privacy.

When configuring QUIC-LB, administrators must evaluate the privacy tradeoff considering the relative value of each of these properties, given the trust model between tenants, the presence of methods to obscure the domain name, and value of address migration in the tenant use cases.

As the plaintext algorithm makes no attempt to conceal the server mapping, these deployments SHOULD simply use a common configuration.

8.5. Stateless Reset Oracle

Section 21.9 of [\[RFC9000\]](#) discusses the Stateless Reset Oracle attack. For a server deployment to be vulnerable, an attacking client must be able to cause two packets with the same Destination CID to arrive at two different servers that share the same cryptographic context for Stateless Reset tokens. As QUIC-LB requires deterministic routing of DCIDs over the life of a connection, it is a sufficient means of avoiding an Oracle without additional measures.

Note also that when a server starts using a new QUIC-LB config rotation codepoint, new CIDs might not be unique with respect to previous configurations that occupied that codepoint, and therefore different clients may have observed the same CID and stateless reset token. A straightforward method of managing stateless reset keys is to maintain a separate key for each config rotation codepoint, and replace each key when the configuration for that codepoint changes. Thus, a server transitions from one config to another, it will be able to generate correct tokens for connections using either type of CID.

8.6. Connection ID Entropy

If a server ever reuses a nonce in generating a CID for a given configuration, it risks exposing sensitive information. Given the same server ID, the CID will be identical (aside from a possible difference in the first octet). This can risk exposure of the QUIC-LB key. If two clients receive the same connection ID, they also have each other's stateless reset token unless that key has changed in the interim.

The encrypt mode needs to generate different cipher text for each generated Connection ID instance to protect the Server ID. To do so, at least four octets of the CID are reserved for a nonce that, if used only once, will result in unique cipher text for each Connection ID.

If servers simply increment the nonce by one with each generated connection ID, then it is safe to use the existing keys until any server's nonce counter exhausts the allocated space and rolls over. To maximize entropy, servers SHOULD start with a random nonce value, in which case the configuration is usable until the nonce value wraps around to zero and then reaches the initial value again.

Whether or not it implements the counter method, the server MUST NOT reuse a nonce until it switches to a configuration with new keys.

If the nonce is sent in plaintext, servers MUST generate nonces so that they appear to be random. Observable correlations between

plaintext nonces would provide trivial linkability between individual connections, rather than just to a common server.

For any algorithm, configuration agents SHOULD implement an out-of-band method to discover when servers are in danger of exhausting their nonce space, and SHOULD respond by issuing a new configuration. A server that has exhausted its nonces MUST either switch to a different configuration, or if none exists, use the 4-tuple routing config rotation codepoint.

When sizing a nonce that is to be randomly generated, the configuration agent SHOULD consider that a server generating a N-bit nonce will create a duplicate about every $2^{(N/2)}$ attempts, and therefore compare the expected rate at which servers will generate CIDs with the lifetime of a configuration.

9. IANA Considerations

There are no IANA requirements.

10. References

10.1. Normative References

[RFC8999] Thomson, M., "Version-Independent Properties of QUIC", RFC 8999, DOI 10.17487/RFC8999, May 2021, <<https://www.rfc-editor.org/info/rfc8999>>.

[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>>.

10.2. Informative References

[I-D.draft-ietf-tls-dtls13] Rescorla, E., Tschofenig, H., and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3", Work in Progress, Internet-Draft, draft-ietf-tls-dtls13-43, 30 April 2021, <<https://www.ietf.org/archive/id/draft-ietf-tls-dtls13-43.txt>>.

[I-D.ietf-tls-dtls-connection-id] Rescorla, E., Tschofenig, H., Fossati, T., and A. Kraus, "Connection Identifier for DTLS 1.2", Work in Progress, Internet-Draft, draft-ietf-tls-dtls-connection-id-13, 22 June 2021, <<https://www.ietf.org/archive/id/draft-ietf-tls-dtls-connection-id-13.txt>>.

[I-D.ietf-tls-esni] Rescorla, E., Oku, K., Sullivan, N., and C. A. Wood, "TLS Encrypted Client Hello", Work in Progress,

Internet-Draft, draft-ietf-tls-esni-14, 13 February 2022, <<https://www.ietf.org/archive/id/draft-ietf-tls-esni-14.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>>.
- [RFC4347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security", RFC 4347, DOI 10.17487/RFC4347, April 2006, <<https://www.rfc-editor.org/info/rfc4347>>.
- [RFC6020] Bjorklund, M., Ed., "YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)", RFC 6020, DOI 10.17487/RFC6020, October 2010, <<https://www.rfc-editor.org/info/rfc6020>>.
- [RFC6347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security Version 1.2", RFC 6347, DOI 10.17487/RFC6347, January 2012, <<https://www.rfc-editor.org/info/rfc6347>>.
- [RFC7696] Housley, R., "Guidelines for Cryptographic Algorithm Agility and Selecting Mandatory-to-Implement Algorithms", BCP 201, RFC 7696, DOI 10.17487/RFC7696, November 2015, <<https://www.rfc-editor.org/info/rfc7696>>.
- [RFC7983] Petit-Huguenin, M. and G. Salgueiro, "Multiplexing Scheme Updates for Secure Real-time Transport Protocol (SRTP) Extension for Datagram Transport Layer Security (DTLS)", RFC 7983, DOI 10.17487/RFC7983, September 2016, <<https://www.rfc-editor.org/info/rfc7983>>.
- [RFC8340] Bjorklund, M. and L. Berger, Ed., "YANG Tree Diagrams", BCP 215, RFC 8340, DOI 10.17487/RFC8340, March 2018, <<https://www.rfc-editor.org/info/rfc8340>>.
- [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>>.

Appendix A. QUIC-LB YANG Model

These YANG models conform to [RFC6020] and express a complete QUIC-LB configuration. There is one model for the server and one for the middlebox (i.e the load balancer and/or Retry Service).

```
module ietf-quic-lb-server {
  yang-version "1.1";
  namespace "urn:ietf:params:xml:ns:yang:ietf-quic-lb";
  prefix "quic-lb";

  import ietf-yang-types {
    prefix yang;
    reference
      "RFC 6991: Common YANG Data Types.";
  }

  import ietf-inet-types {
    prefix inet;
    reference
      "RFC 6991: Common YANG Data Types.";
  }

  organization
    "IETF QUIC Working Group";

  contact
    "WG Web: <http://datatracker.ietf.org/wg/quic>
    WG List: <quic@ietf.org>

    Authors: Martin Duke (martin.h.duke at gmail dot com)
            Nick Banks (nibanks at microsoft dot com)
            Christian Huitema (huitema at huitema.net);

  description
    "This module enables the explicit cooperation of QUIC servers with
    trusted intermediaries without breaking important protocol
    features.

    Copyright (c) 2022 IETF Trust and the persons identified as
    authors of the code. All rights reserved.

    Redistribution and use in source and binary forms, with or
    without modification, is permitted pursuant to, and subject to
    the license terms contained in, the Simplified BSD License set
    forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX
    (https://www.rfc-editor.org/info/rfcXXXX); see the RFC itself
    for full legal notices.

    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 (RFC 2119) (RFC 8174) when, and only when, they appear in all capitals, as shown here.";

```
revision "2022-02-11" {
  description
    "Updated to design in version 13 of the draft";
  reference
    "RFC XXXX, QUIC-LB: Generating Routable QUIC Connection IDs";
}

container quic-lb {
  presence "The container for QUIC-LB configuration.";

  description
    "QUIC-LB container.";

  typedef quic-lb-key {
    type yang:hex-string {
      length 47;
    }
    description
      "This is a 16-byte key, represented with 47 bytes";
  }

  leaf config-id {
    type uint8 {
      range "0..2";
    }
    mandatory true;
    description
      "Identifier for this CID configuration.";
  }

  leaf first-octet-encodes-cid-length {
    type boolean;
    default false;
    description
      "If true, the six least significant bits of the first CID
      octet encode the CID length minus one.";
  }

  leaf server-id-length {
    type uint8 {
      range "1..15";
    }
    must '. <= (19 - ../nonce-length)' {
      error-message
        "Server ID and nonce lengths must sum to no more than 19.";
    }
  }
}
```



```
    mandatory true;
    description
      "Length (in octets) of a server ID. Further range-limited
        by nonce-length.";
  }

  leaf nonce-length {
    type uint8 {
      range "4..18";
    }
    mandatory true;
    description
      "Length, in octets, of the nonce. Short nonces mean there will
        be frequent configuration updates.";
  }

  leaf cid-key {
    type quic-lb-key;
    description
      "Key for encrypting the connection ID.";
  }

  leaf server-id {
    type yang:hex-string;
    must "string-length(.) = 3 * ../../server-id-length - 1";
    mandatory true;
    description
      "An allocated server ID";
  }
}
}
```

```
module ietf-quic-lb-middlebox {
  yang-version "1.1";
  namespace "urn:ietf:params:xml:ns:yang:ietf-quic-lb";
  prefix "quic-lb";

  import ietf-yang-types {
    prefix yang;
    reference
      "RFC 6991: Common YANG Data Types.";
  }

  import ietf-inet-types {
    prefix inet;
    reference
      "RFC 6991: Common YANG Data Types.";
  }

  organization
    "IETF QUIC Working Group";

  contact
    "WG Web: <http://datatracker.ietf.org/wg/quic>
    WG List: <quic@ietf.org>

    Authors: Martin Duke (martin.h.duke at gmail dot com)
            Nick Banks (nibanks at microsoft dot com)
            Christian Huitema (huitema at huitema.net);

  description
    "This module enables the explicit cooperation of QUIC servers with
    trusted intermediaries without breaking important protocol
    features.

    Copyright (c) 2021 IETF Trust and the persons identified as
    authors of the code. All rights reserved.

    Redistribution and use in source and binary forms, with or
    without modification, is permitted pursuant to, and subject to
    the license terms contained in, the Simplified BSD License set
    forth in Section 4.c of the IETF Trust's Legal Provisions
    Relating to IETF Documents
    (https://trustee.ietf.org/license-info).

    This version of this YANG module is part of RFC XXXX
    (https://www.rfc-editor.org/info/rfcXXXX); see the RFC itself
    for full legal notices.

    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 (RFC 2119) (RFC 8174) when, and only when, they appear in all capitals, as shown here.";

```
revision "2021-02-11" {
  description
    "Updated to design in version 13 of the draft";
  reference
    "RFC XXXX, QUIC-LB: Generating Routable QUIC Connection IDs";
}

container quic-lb {
  presence "The container for QUIC-LB configuration.";

  description
    "QUIC-LB container.";

  typedef quic-lb-key {
    type yang:hex-string {
      length 47;
    }
    description
      "This is a 16-byte key, represented with 47 bytes";
  }

  list cid-configs {
    key "config-rotation-bits";
    description
      "List up to three load balancer configurations";

    leaf config-rotation-bits {
      type uint8 {
        range "0..2";
      }
      mandatory true;
      description
        "Identifier for this CID configuration.";
    }

    leaf server-id-length {
      type uint8 {
        range "1..15";
      }
      must '. <= (19 - ../nonce-length)' {
        error-message
          "Server ID and nonce lengths must sum to no more than 19.";
      }
      mandatory true;
      description
        "Length (in octets) of a server ID. Further range-limited
```



```

module: ietf-quic-lb-server
  +--rw quic-lb!
    +--rw config-id                               uint8
    +--rw first-octet-encodes-cid-length?         boolean
    +--rw server-id-length                       uint8
    +--rw nonce-length                           uint8
    +--rw cid-key?                               quic-lb-key
    +--rw server-id                              yang:hex-string

```

```

module: ietf-quic-lb-middlebox
  +--rw quic-lb!
    +--rw cid-configs* [config-rotation-bits]
      | +--rw config-rotation-bits   uint8
      | +--rw server-id-length       uint8
      | +--rw cid-key?               quic-lb-key
      | +--rw nonce-length           uint8
      | +--rw server-id-mappings* [server-id]
      |   +--rw server-id            yang:hex-string
      |   +--rw server-address       inet:ip-address

```

Appendix B. Load Balancer Test Vectors

This section uses the following abbreviations:

cid	Connection ID
cr_bits	Config Rotation Bits
LB	Load Balancer
sid	Server ID

In all cases, the server is configured to encode the CID length.

B.1. Unencrypted CIDs

```

cr_bits sid nonce cid
0 c4605e 4504cc4f 07c4605e4504cc4f
1 350d28b420 3487d970b 40a350d28b4203487d970b

```

B.2. Encrypted CIDs

The key for all of these examples is 8f95f09245765f80256934e50c66207f. The test vectors include an example that uses the 16-octet single-pass special case, as well as an instance where the server ID length exceeds the nonce length, requiring a fourth decryption pass.

```
cr_bits sid nonce cid
0 ed793a ee080dbf 07fbfe05f731b425
1 ed793a51d49b8f5fab65 ee080dbf48 4f010956fb5c1d4d86e010183e0b7d1e
2 ed793a51d49b8f5f ee080dbf48c0d1e5 904dd2d05a7b0de9b2b9907afb5ecf8cc3
0 ed793a51d49b8f5fab ee080dbf48c0d1e55d 127a285a09f85280f4fd6abb434a7159e4d3eb
```

Appendix C. Interoperability with DTLS over UDP

Some environments may contain DTLS traffic as well as QUIC operating over UDP, which may be hard to distinguish.

In most cases, the packet parsing rules above will cause a QUIC-LB load balancer to route DTLS traffic in an appropriate way. DTLS 1.3 implementations that use the `connection_id` extension [[I-D.ietf-tls-dtls-connection-id](#)] might use the techniques in this document to generate connection IDs and achieve robust routability for DTLS associations if they meet a few additional requirements. This non-normative appendix describes this interaction.

C.1. DTLS 1.0 and 1.2

DTLS 1.0 [[RFC4347](#)] and 1.2 [[RFC6347](#)] use packet formats that a QUIC-LB router will interpret as short header packets with CIDs that request 4-tuple routing. As such, they will route such packets consistently as long as the 4-tuple does not change. Note that DTLS 1.0 has been deprecated by the IETF.

The first octet of every DTLS 1.0 or 1.2 datagram contains the content type. A QUIC-LB load balancer will interpret any content type less than 128 as a short header packet, meaning that the subsequent octets should contain a connection ID.

Existing TLS content types comfortably fit in the range below 128. Assignment of codepoints greater than 64 would require coordination in accordance with [[RFC7983](#)], and anyway would likely create problems demultiplexing DTLS and version 1 of QUIC. Therefore, this document believes it is extremely unlikely that TLS content types of 128 or greater will be assigned. Nevertheless, such an assignment would cause a QUIC-LB load balancer to interpret the packet as a QUIC long header with an essentially random connection ID, which is likely to be routed irregularly.

The second octet of every DTLS 1.0 or 1.2 datagram is the bitwise complement of the DTLS Major version (i.e. version 1.x = 0xfe). A QUIC-LB load balancer will interpret this as a connection ID that requires 4-tuple based load balancing, meaning that the routing will be consistent as long as the 4-tuple remains the same.

[[I-D.ietf-tls-dtls-connection-id](#)] defines an extension to add connection IDs to DTLS 1.2. Unfortunately, a QUIC-LB load balancer will not correctly parse the connection ID and will continue 4-tuple routing. An modified QUIC-LB load balancer that correctly identifies DTLS and parses a DTLS 1.2 datagram for the connection ID is outside the scope of this document.

C.2. DTLS 1.3

DTLS 1.3 [[I-D.draft-ietf-tls-dtls13](#)] changes the structure of datagram headers in relevant ways.

Handshake packets continue to have a TLS content type in the first octet and 0xfe in the second octet, so they will be 4-tuple routed, which should not present problems for likely NAT rebinding or address change events.

Non-handshake packets always have zero in their most significant bit and will therefore always be treated as QUIC short headers. If the connection ID is present, it follows in the succeeding octets. Therefore, a DTLS 1.3 association where the server utilizes Connection IDs and the encodings in this document will be routed correctly in the presence of client address and port changes.

However, if the client does not include the `connection_id` extension in its ClientHello, the server is unable to use connection IDs. In this case, non-handshake packets will appear to contain random connection IDs and be routed randomly. Thus, unmodified QUIC-LB load balancers will not work with DTLS 1.3 if the client does not advertise support for connection IDs, or the server does not request the use of a compliant connection ID.

A QUIC-LB load balancer might be modified to identify DTLS 1.3 packets and correctly parse the fields to identify when there is no connection ID and revert to 4-tuple routing, removing the server requirement above. However, such a modification is outside the scope of this document, and classifying some packets as DTLS might be incompatible with future versions of QUIC.

C.3. Future Versions of DTLS

As DTLS does not have an IETF consensus document that defines what parts of DTLS will be invariant in future versions, it is difficult to speculate about the applicability of this section to future versions of DTLS.

Appendix D. Acknowledgments

Manasi Deval, Erik Fuller, Toma Gavrichenkov, Jana Iyengar, Subodh Iyengar, Ladislav Lhotka, Jan Lindblad, Ling Tao Nju, Ilari

Liusvaara, Kazuho Oku, Udip Pant, Ian Swett, Martin Thomson, Dmitri Tikhonov, Victor Vasiliev, and William Zeng Ke all provided useful input to this document.

Appendix E. Change Log

RFC Editor's Note: Please remove this section prior to publication of a final version of this document.

E.1. since draft-ietf-quic-load-balancers-12

*Separated Retry Service design into a separate draft.

E.2. since draft-ietf-quic-load-balancers-11

*Fixed mistakes in test vectors

E.3. since draft-ietf-quic-load-balancers-10

*Refactored algorithm descriptions; made the 4-pass algorithm easier to implement

*Revised test vectors

*Split YANG model into a server and middlebox version

E.4. since draft-ietf-quic-load-balancers-09

*Renamed "Stream Cipher" and "Block Cipher" to "Encrypted Short" and "Encrypted Long"

*Added section on per-connection state

*Changed "Encrypted Short" to a 4-pass algorithm.

*Recommended a random initial nonce when incrementing.

*Clarified what SNI LBs should do with unknown QUIC versions.

E.5. since draft-ietf-quic-load-balancers-08

*Eliminate Dynamic SID allocation

*Eliminated server use bytes

E.6. since draft-ietf-quic-load-balancers-07

*Shortened SSCID nonce minimum length to 4 bytes

*Removed RSCID from Retry token body

- *Simplified CID formats

- *Shrunk size of SID table

E.7. since draft-ietf-quic-load-balancers-06

- *Added interoperability with DTLS

- *Changed "non-compliant" to "unroutable"

- *Changed "arbitrary" algorithm to "fallback"

- *Revised security considerations for mistrustful tenants

- *Added retry service considerations for non-Initial packets

E.8. since draft-ietf-quic-load-balancers-05

- *Added low-config CID for further discussion

- *Complete revision of shared-state Retry Token

- *Added YANG model

- *Updated configuration limits to ensure CID entropy

- *Switched to notation from quic-transport

E.9. since draft-ietf-quic-load-balancers-04

- *Rearranged the shared-state retry token to simplify token processing

- *More compact timestamp in shared-state retry token

- *Revised server requirements for shared-state retries

- *Eliminated zero padding from the test vectors

- *Added server use bytes to the test vectors

- *Additional compliant DCID criteria

E.10. since-draft-ietf-quic-load-balancers-03

- *Improved Config Rotation text

- *Added stream cipher test vectors

- *Deleted the Obfuscated CID algorithm

E.11. since-draft-ietf-quic-load-balancers-02

- *Replaced stream cipher algorithm with three-pass version
- *Updated Retry format to encode info for required TPs
- *Added discussion of version invariance
- *Cleaned up text about config rotation
- *Added Reset Oracle and limited configuration considerations
- *Allow dropped long-header packets for known QUIC versions

E.12. since-draft-ietf-quic-load-balancers-01

- *Test vectors for load balancer decoding
- *Deleted remnants of in-band protocol
- *Light edit of Retry Services section
- *Discussed load balancer chains

E.13. since-draft-ietf-quic-load-balancers-00

- *Removed in-band protocol from the document

E.14. Since draft-duke-quic-load-balancers-06

- *Switch to IETF WG draft.

E.15. Since draft-duke-quic-load-balancers-05

- *Editorial changes
- *Made load balancer behavior independent of QUIC version
- *Got rid of token in stream cipher encoding, because server might not have it
- *Defined "non-compliant DCID" and specified rules for handling them.
- *Added psuedocode for config schema

E.16. Since draft-duke-quic-load-balancers-04

- *Added standard for retry services

E.17. Since draft-duke-quic-load-balancers-03

- *Renamed Plaintext CID algorithm as Obfuscated CID
- *Added new Plaintext CID algorithm
- *Updated to allow 20B CIDs
- *Added self-encoding of CID length

E.18. Since draft-duke-quic-load-balancers-02

- *Added Config Rotation
- *Added failover mode
- *Tweaks to existing CID algorithms
- *Added Block Cipher CID algorithm
- *Reformatted QUIC-LB packets

E.19. Since draft-duke-quic-load-balancers-01

- *Complete rewrite
- *Supports multiple security levels
- *Lightweight messages

E.20. Since draft-duke-quic-load-balancers-00

- *Converted to markdown
- *Added variable length connection IDs

Authors' Addresses

Martin Duke
Google

Email: martin.h.duke@gmail.com

Nick Banks
Microsoft

Email: nibanks@microsoft.com

Christian Huitema
Private Octopus Inc.

Email: huitema@huitema.net