

QUIC
Internet-Draft
Intended status: Standards Track
Expires: 28 February 2026

M. Duke
Google
N. Banks
Microsoft
C. Huitema
Private Octopus Inc.
27 August 2025

QUIC-LB: Generating Routable QUIC Connection IDs
draft-ietf-quic-load-balancers-21

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.

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 28 February 2026.

Copyright Notice

Copyright (c) 2025 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	4
1.1. Terminology	5
1.2. Notation	5
2. Overview	5
3. First CID octet	6
3.1. Config Rotation	6
3.2. Configuration Failover	7
3.3. Length Self-Description	8
3.4. Format	8
4. Unroutable Connection IDs	8
4.1. Definition	8
4.2. Load Balancer Forwarding	9
4.3. Fallback Algorithms	10
4.3.1. Baseline Fallback Algorithm	10
4.3.2. Advanced Fallback Algorithm	11
5. Server ID Encoding in Connection IDs	11
5.1. Server ID Allocation	12
5.2. CID format	12
5.3. Configuration Agent Actions	13
5.4. Server Actions	13
5.4.1. Special Case: Single Pass Encryption	14
5.4.2. General Case: Four-Pass Encryption	14
5.5. Load Balancer Actions	19
5.5.1. Special Case: Single Pass Encryption	19
5.5.2. General Case: Four-Pass Encryption	19
6. Per-connection state	20
7. Additional Use Cases	21
7.1. Load balancer chains	21
7.2. Server Process Demultiplexing	22
7.3. Moving connections between servers	22
8. Version Invariance of QUIC-LB	23
9. Security Considerations	24
9.1. Attackers not between the load balancer and server	25
9.2. Attackers between the load balancer and server	25
9.3. Multiple Configuration IDs	25
9.4. Limited configuration scope	25
9.5. Stateless Reset Oracle	26

9.6. Connection ID Entropy	27
9.7. Distinguishing Attacks	28
9.8. Early deletion of load balancer connection state	28
10. IANA Considerations	29
11. References	29
11.1. Normative References	29
11.2. Informative References	29
Appendix A. QUIC-LB YANG Model	30
A.1. Tree Diagram	36
Appendix B. Load Balancer Test Vectors	37
B.1. Unencrypted CIDs	37
B.2. Encrypted CIDs	37
Appendix C. Interoperability with DTLS over UDP	38
C.1. DTLS 1.0 and 1.2	38
C.2. DTLS 1.3	39
C.3. Future Versions of DTLS	39
Appendix D. Acknowledgments	39
Appendix E. Change Log	40
E.1. since draft-ietf-quic-load-balancers-20	40
E.2. since draft-ietf-quic-load-balancers-19	40
E.3. since draft-ietf-quic-load-balancers-18	40
E.4. since draft-ietf-quic-load-balancers-17	40
E.5. since draft-ietf-quic-load-balancers-16	40
E.6. since draft-ietf-quic-load-balancers-15	40
E.7. since draft-ietf-quic-load-balancers-14	40
E.8. since draft-ietf-quic-load-balancers-13	41
E.9. since draft-ietf-quic-load-balancers-12	41
E.10. since draft-ietf-quic-load-balancers-11	41
E.11. since draft-ietf-quic-load-balancers-10	41
E.12. since draft-ietf-quic-load-balancers-09	41
E.13. since draft-ietf-quic-load-balancers-08	41
E.14. since draft-ietf-quic-load-balancers-07	42
E.15. since draft-ietf-quic-load-balancers-06	42
E.16. since draft-ietf-quic-load-balancers-05	42
E.17. since draft-ietf-quic-load-balancers-04	42
E.18. since draft-ietf-quic-load-balancers-03	43
E.19. since draft-ietf-quic-load-balancers-02	43
E.20. since draft-ietf-quic-load-balancers-01	43
E.21. since draft-ietf-quic-load-balancers-00	43
E.22. Since draft-duke-quic-load-balancers-06	43
E.23. Since draft-duke-quic-load-balancers-05	43
E.24. Since draft-duke-quic-load-balancers-04	44
E.25. Since draft-duke-quic-load-balancers-03	44
E.26. Since draft-duke-quic-load-balancers-02	44
E.27. Since draft-duke-quic-load-balancers-01	44
E.28. Since draft-duke-quic-load-balancers-00	44
Authors' Addresses	44

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 encode useful routing information for load balancers in connection IDs. 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].

2. Overview

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 specification 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.

[RFC8999] specifies that endpoints generate their own connection IDs, implying that all QUIC versions will have a mechanism to communicate their connection IDs to the peer. In QUIC version 1 and 2, the server does so using the Source Connection ID field of its long header packets for the first connection ID, and NEW_CONNECTION_ID frames for subsequent CIDs.

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 three bits of the connection ID to multiplex incoming Destination Connection IDs (DCIDs) over these schemes (see Section 3.1).

3. First CID octet

The Connection ID construction schemes defined in this document reserve the first octet of a CID 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."

3.1. Config Rotation

The first three 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). A configuration MUST NOT use the reserved identifier 0b111 (see Section 3.2 below).

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 use that QUIC version's methods to update the client with CIDs (e.g., NEW_CONNECTION_ID frames) using the new configuration and retire CIDs using the old configuration.

It is also possible to use these bits for more long-lived distinction of different configurations, but this has privacy implications (see Section 9.3).

3.2. Configuration Failover

In some deployments, an infrastructure will not receive traffic unless all servers have received a configuration, and load balancers have a superset of all configurations that are active in the server pool, thus guaranteeing that any CID generated by a server is decodable by any load balancer. Servers and load balancers deployed under all of these assumptions can ignore the provisions in this subsection.

Load balancers treat connection IDs for which they have no corresponding config ID as unroutable (see Section 4). If they have no configuration at all, then all connection IDs are unroutable.

Servers with no active configuration MUST issue connection IDs with the reserved value of the three most significant bits set to 0b111 to signify the connection ID is unroutable. These connection IDs MUST self-encode their length (see Section 3.3).

Servers with no active configuration SHOULD provide the client exactly one CID over the life of the connection. In QUIC versions 1 and 2, therefore, servers SHOULD NOT send any NEW_CONNECTION_ID frames, instead delivering a single CID via the Source Connection ID of long headers it sends.

Servers with no active configuration SHOULD send the "disable_active_migration" transport parameter, or a similar message in future QUIC versions.

When using codepoint 0b111, all bytes but the first SHOULD have no larger of a chance of collision as random bytes. The connection ID SHOULD be of at least length 8 to provide 7 bytes of entropy after the first octet with a low chance of collision.

3.3. 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 might be impractical to efficiently lookup keys if the connection ID varies in length and 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 5 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 choose the five bits so as to have no observable relationship to previous connection IDs issued for that connection.

3.4. Format

```
First Octet {  
    Config Rotation (3),  
    CID Len or Random Bits (5),  
}
```

Figure 1: 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.

4. Unroutable Connection IDs

4.1. Definition

QUIC-LB servers with a valid configuration 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.

Furthermore, servers without a valid configuration, or a configuration not present at the load balancer, will also generate connection IDs that are not decodable, and these CIDs are likely to persist for the duration of the connection.

These 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 3.1) do not correspond to an active configuration. Note: a packet with a DCID with config ID codepoint 0b111 (see Section 3.2) is always unroutable.
- * If the packet header encodes the DCID length, the DCID is not long enough for the decoder to process.
- * The extracted server mapping does not correspond to an active server.

If the load balancer has knowledge that all servers in the pool are encoding CID length in the first octet (see Section 3.3), it MAY perform additional checks based on that self-encoded length:

- * In a long header, verify that the self-encoded length is consistent with the CID length field in the header (i.e. the self-encoded length is one less)
- * Verify that the self-encoded length is consistent with the QUIC version, if known.
- * Verify that the self-encoded length is large enough for the decoder to process using the indicated config ID.

DCIDs that do not meet any of these criteria are routable.

4.2. Load Balancer Forwarding

Load balancers execute the following steps in order until one results in a routing decision. The steps refer to state that some load balancers will maintain, depending on the deployment's underlying assumptions. See Section 4.3 for further discussion of this state.

1. If the packet contains a routable CID, route the packet accordingly.

2. If the packet has a long header and matches an entry in a table of routing decisions indexed by a concatenation of 4-tuple and Source CID, route the packet accordingly.
3. If the packet matches an entry in a table of routing decisions by destination CID, route the packet accordingly.
4. If packet matches an entry in a table of routing decisions by 4-tuple, route the packet accordingly.
5. Use the fallback algorithm to make a routing decision and, if applicable, record the results in the tables indexed by 4-tuple and/or CID. In some cases, described below, the load balancer might buffer the packet to defer a decision.

4.3. Fallback Algorithms

There are conditions described above where a load balancer routes a packet using a "fallback algorithm." A standardized algorithm design is not necessary for interoperability, so load balancers can implement any algorithm that meets the relevant requirements below.

There is a baseline case that has relatively simple requirements of the chosen fallback algorithm, and an advanced case with more capabilities and more complex requirements.

4.3.1. Baseline Fallback Algorithm

All load balancers **MUST** implement a baseline fallback algorithm that takes only the 4-tuple as an input and outputs a routing decision.

If it is impossible for the server to generate CIDs that the load balancer cannot decode (see Section 3.2), there are no further requirements in this subsection.

Otherwise, the load balancer **SHOULD** maintain a table of 4-tuples that carried unroutable DCIDs and the resulting routing decision. Provided the table does not overflow, and the load balancer does not lose state, this allows connections to survive when the server pool changes, which would sometimes change the output of the fallback algorithm.

The load balancer **MAY** maintain a table of observed unroutable DCIDs and the resulting routing decision. Provided the table does not overflow, these connections will be robust to NAT rebinding.

Load balancers **SHOULD** maintain per-flow timers to periodically purge state in the tables described above.

4.3.2. Advanced Fallback Algorithm

Some architectures might require a load balancer to choose a server pool based on deep packet inspection of a client packet. For example, it may use the TLS 1.3 Server Name Indication (SNI) ([RFC6066]) field. The advanced fallback algorithm enables this capability but levies several additional requirements to make consistent routing decisions.

For packets not known to belong to a QUIC version the load balancer can parse, load balancers **MUST** use the baseline fallback algorithm if the DCID is unroutable.

For known QUIC versions, the fallback algorithm **MAY** parse packets and use that information to make a routing decision.

If so, it **MUST** have the ability to buffer packets with unroutable DCIDs to await further packets that allow it to make a routing decision, as the fields of interest can be an arbitrary number of packets into the connection.

4-tuple routing is not sufficient for this use case, because a client can use the same 4-tuple for two connections that should be routed differently (e.g. because they target different SNIs), as long as the packet contains a source connection ID of nonzero length.

Therefore, the load balancer **SHOULD** maintain two tables that map different values to a routing decision:

- * a table indexed by a concatenation of the 4-tuple and source CID, which might be zero-length, to route subsequent long header packets that do not contain the server-generated connection ID;
- * a table indexed by destination CID, if and only if it is possible for the server to generate unroutable CIDs. This table can be shared with the one in use for the baseline fallback algorithm.

If either table overflows, or if the load balancer loses state, it is likely the load balancer will misroute packets.

Load balancers **SHOULD** maintain per-flow timers to periodically purge state in the tables described above.

5. Server ID Encoding in Connection IDs

5.1. 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 7.1).

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.

5.2. CID format

All connection IDs use the following format:

```
QUIC-LB Connection ID {
    First Octet (8),
    Plaintext Block (40..152),
}
Plaintext Block {
    Server ID (8..),
    Nonce (32..),
}
```

Figure 2: CID Format

The First Octet field serves one or two purposes, as defined in Section 3.

The Server ID field encodes the information necessary for the load balancer to route a packet with that connection ID. It is often encrypted.

The server uses the Nonce field to make sure that each connection ID it generates is unique, even though they all use the same Server ID.

5.3. Configuration Agent Actions

The configuration agent assigns a server ID to every server in its pool in accordance with Section 5.1, 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.

5.4. 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 have no observable relationship to the field in previously issued connection IDs. If there is a key, the server fills the nonce field with a nonce of its choosing. See Section 9.6 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 NOT** provide any observable correlation to previous connection IDs for that connection (e.g., the bytes can be chosen at random).

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 [NIST-AES-ECB]. 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.

5.4.1. 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.

5.4.2. General Case: Four-Pass Encryption

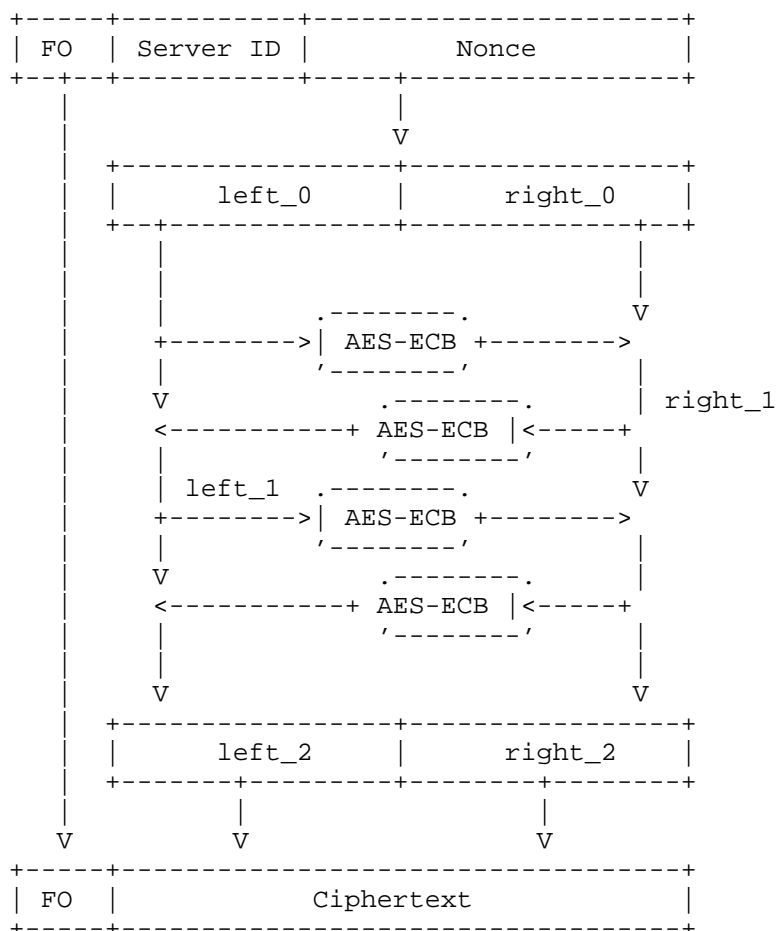
Any other field length requires four passes for encryption and at least three for decryption. To understand this algorithm, it is useful to define four functions that minimize the amount of bit-shifting necessary in the event that there are an odd number of octets.

When configured with both a key, and a nonce length and server ID length that sum to any number other than 16, the server MUST follow the algorithm below to encrypt the connection ID.

5.4.2.1. Overview

The 4-pass algorithm is a four-round Feistel Network with the round function being AES-ECB. Most modern applications of Feistel Networks have more than four rounds. The implications of this choice, which is meant to limit the per-packet compute overhead at load balancers, are discussed in Section 9.7.

The server concatenates the server ID and nonce into a single field, which is then split into equal halves. In successive passes, one of these halves is expanded into a 16B plaintext, encrypted with AES-ECB, and the result XORed with the other half. The diagram below shows the conceptual processing of a plaintext server ID and nonce into a connection ID. 'FO' stands for 'First Octet'.



5.4.2.2. Useful functions

Two functions are useful to define:

The `expand(length, pass, input_bytes)` function concatenates three arguments and outputs 16 zero-padded octets.

The output of `expand` is as follows:

```
ExpandResult {
    input_bytes(...),
    ZeroPad(...),
    length(8),
    pass(8)
}
```

in which:

- * 'input_bytes' is drawn from one half of the plaintext. It forms the N most significant octets of the output, where N is half the 'length' argument, rounded up, and thus a number between 3 and 10, inclusive.
- * 'Zeropad' is a set of 14-N octets set to zero.
- * 'length' is an 8-bit integer that reports the sum of the configured nonce length and server id length in octets, and forms the fifteenth octet of the output. The 'length' argument MUST NOT exceed 19 and MUST NOT be less than 5.
- * 'pass' is an 8-bit integer that reports the 'pass' argument of the algorithm, and forms the sixteenth (least significant) octet of the output. It guarantees that the cryptographic input of every pass of the algorithm is unique.

For example,

```
expand(0x06, 0x02, 0xaaba3c) = 0xaaba3c000000000000000000000000602
```

Similarly, `truncate(input, n)` returns the first n octets of 'input'.

```
truncate(0x2094842ca49256198c2deaa0ba53caa0, 4) = 0x2094842c
```

Let 'half_len' be equal to 'plaintext_len' / 2, rounded up.

5.4.2.3. Algorithm Description

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`. The length of the result in octets is `plaintext_len`.
2. The server splits `plaintext_CID` into components `left_0` and `right_0` of equal length `half_len`. If `plaintext_len` is odd, `right_0` clears its first four bits, and `left_0` clears its last four bits. For example, `0x7040b81b55ccf3` would split into a `left_0` of `0x7040b810` and `right_0` of `0x0b55ccf3`.
3. Encrypt the result of `expand(plaintext_len, 1, left_0)` using an AES-ECB-128 cipher to obtain a ciphertext.

4. XOR the first `half_len` octets of the ciphertext with `right_0` to form `right_1`. Steps 3 and 4 can be summarized as

```
result = AES_ECB(key, expand(plaintext_len, 1, left_0))
right_1 = XOR(right_0, truncate(result, half_len))
```

5. If the `plaintext_len` is odd, clear the first four bits of `right_1`.
6. Repeat steps 3 and 4, but use them to compute `left_1` by expanding and encrypting `right_1` with `pass = 2`, and XOR the results with `left_0`.

```
result = AES_ECB(key, expand(plaintext_len, 2, right_1))
left_1 = XOR(left_0, truncate(result, half_len))
```

7. If the `plaintext_len` is odd, clear the last four bits of `left_1`.
8. Repeat steps 3 and 4, but use them to compute `right_2` by expanding and encrypting `left_1` with `pass = 3`, and XOR the results with `right_1`.

```
result = AES_ECB(key, expand(plaintext_len, 3, left_1))
right_2 = XOR(right_1, truncate(result, half_len))
```

9. If the `plaintext_len` is odd, clear the first four bits of `right_2`.
10. Repeat steps 3 and 4, but use them to compute `left_2` by expanding and encrypting `right_2` with `pass = 4`, and XOR the results with `left_1`.

```
result = AES_ECB(key, expand(plaintext_len, 4, right_2))
left_2 = XOR(left_1, truncate(result, half_len))
```

11. If the `plaintext_len` is odd, clear the last four bits of `left_2`.
12. The server concatenates `left_2` with `right_2` to form the ciphertext CID, which it appends to the first octet. If `plaintext_len` is odd, the four least significant bits of `left_2` and four most significant bits of `right_2`, which are all zero, are stripped off before concatenation to make the resulting ciphertext the same length as the original plaintext.

5.4.2.4. Encryption Example

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.

```
server_id = 0x31441a
nonce = 0x9c69c275
key = 0xfdf726a9893ec05c0632d3956680baf0

// step 1
plaintext_CID = 0x31441a9c69c275
plaintext_len = 7

// step 2
hash_len = 4
left_0 = 0x31441a90
right_0 = 0x0c69c275

// step 3
aes_input = 0x31441a90000000000000000000000000701
aes_output = 0xa255dd8cdacf01948d3a848c3c7fee23

// step 4
right_1 = 0x0c69c275 ^ 0xa255dd8c = 0xae3c1ff9

// step 5 (clear bits)
right_1 = 0x0e3c1ff9

// step 6
aes_input = 0x0e3c1ff9000000000000000000000000702
aes_output = 0xe5e452cb9e1bedb0b2bf830506bf4c4e
left_1 = 0x31441a90 ^ 0xe5e452cb = 0xd4a0485b

// step 7 (clear bits)
left_1 = 0xd4a04850

// step 8
aes_input = 0xd4a04850000000000000000000000000703
aes_output = 0xb7821ab3024fed0913b6a04d18e3216f
right_2 = 0x0e3c1ff9 ^ 0xb7821ab3 = 0xb9be054a

// step 9 (clear bits)
right_2 = 0x09be054a

// step 10
aes_input = 0x09be054a000000000000000000000000704
aes_output = 0xb334357cfd81e3f81e180154eaf7378
```

```
left_2 = 0xd4a04850 ^ 0xb3e4357c = 0x67947d2c

// step 11 (clear bits)
left_2 = 0x67947d20

// step 12
cid = first_octet || left_2 || right_2 = 0x0767947d29be054a
```

5.5. Load Balancer Actions

On each incoming packet, the load balancer extracts consecutive octets, beginning with the second octet. If there is no key, the first octets correspond to the server ID.

If there is a key, the load balancer takes one of two actions:

5.5.1. Special Case: Single Pass Encryption

If server ID length and nonce length sum to exactly 16 octets, they form a ciphertext block. The load balancer decrypts the block using the AES-ECB key and extracts the server ID from the most significant bytes of the resulting plaintext.

5.5.2. 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:

```
result = AES_ECB(key, expand(plaintext_len, 4, right_2))
left_1 = XOR(left_2, truncate(result, half_len))
if (plaintext_len_is_odd()) clear_last_bits(left_1, 4)

result = AES_ECB(key, expand(plaintext_len, 3, left_1))
right_1 = XOR(right_2, truncate(result, half_len))
if (plaintext_len_is_odd()) clear_first_bits(left_1, 4)

result = AES_ECB(key, expand(plaintext_len, 2, right_1))
left_0 = XOR(left_1, truncate(result, half_len))
if (plaintext_len_is_odd()) clear_last_bits(left_0, 4)
```

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:

```
result = AES_ECB(key, expand(plaintext_len, 1, left_0))
right_0 = XOR(right_1, truncate(result, half_len))
if (plaintext_len_is_odd()) clear_first_bits(right_0, 4)
```

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

6. Per-connection state

The CID allocation methods QUIC-LB defines no per-connection state at the load balancer, with a few conditional exceptions described in Section 4. Otherwise, the load balancer can extract the server ID from the connection ID of each incoming packet and route that packet accordingly.

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

- * The load balancer only extracts the server ID once until the 4-tuple or connection ID changes. When the CID is encrypted, this might reduce 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 packets that 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.

Note that some heuristics to purge state early can introduce Denial of Service vulnerabilities. For example, one heuristic might delete flow state once the load balancer observes a routable CID on that flow. An attacker that can observe a target flow can store a routable CID from a previous connection and spoof the target flow's 4-tuple with the routable CID, causing premature deletion of that state.

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.

7. Additional Use Cases

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

7.1. 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. Load balancers closer to the client can then treat any parts of the server ID they did not use as part of the nonce.

7.2. Server Process Demultiplexing

QUIC servers might have QUIC running on multiple processes or threads listening on the same address, and have a need to demultiplex between them. In principle, this demultiplexer is a Layer 4 load balancer, and the guidance in Section 7.1 applies. However, in many deployments the demultiplexer lacks the capability to perform decryption operations. Internal server coordination is out of scope of this specification, but this non-normative section proposes some approaches that could work given certain server capabilities:

- * Some bytes of the server ID are reserved to encode the process ID. The demultiplexer might operate based on the 4-tuple or other legacy indicator, but the receiving server process extracts the server ID, and if it does not match the one for that process, the process could "toss" the packet to the correct destination process.
- * Each process could register the connection IDs it generates with the demultiplexer, which routes those connection IDs accordingly.
- * In a combination of the two approaches above, the demultiplexer generally routes by 4-tuple. After a migration, the process tosses the first flight of packets and registers the new connection ID with the demultiplexer. This alternative limits the bandwidth consumption of tossing and the memory footprint of a full connection ID table.
- * When generating a connection ID, the server writes the process ID to the random field of the first octet, or if this is being used for length encoding, in an octet it appends after the ciphertext. It then applies a keyed hash (with a key locally generated for the sole use of that server). The hash result is used as a bitmask to XOR with the bits encoding the process ID. On packet receipt, the demultiplexer applies the same keyed hash to generate the same mask and recover the process ID. (Note that this approach is conceptually similar to QUIC header protection). It is important that the server also appends the process ID to the server ID in the plaintext, so that different processes do not generate the same ciphertext. The load balancer will consider this data to be part of the nonce.

7.3. 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 old server using the "Retire Prior To" field in that frame.

8. 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.

A QUIC version might define limits on connection ID length that make some or all of the mechanisms in this document unusable. For example, a maximum connection ID length could be below the minimum necessary to use all or part of this specification; or, the minimum connection ID length could be larger than the largest value in this specification. Similarly, the length self-encoding specification cannot accommodate connection IDs longer than 32 bytes.

The advanced fallback implementation supports a requirement to inspect version- specific elements of packets to make a routing decision, such as the Server Name Indication (SNI) extension in the TLS Client Hello. 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 version-dependent. Such a load balancer, when it receives packets from an unknown QUIC version, might misdirect initial packets to the wrong tenant. While this can be inefficient, the design in this document 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.

Section 4.2 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 about packets that contain client-generated DCIDs that are not specified in RFC 8999:

1. they do not have short headers;
2. the 4-tuple remains constant;
3. if the load-balancer uses the Advanced Fallback Algorithm, the packets have a constant Source Connection ID.

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 not valid, this specification is likely to lead to loss of packets that contain unroutable DCIDs, and in extreme cases connection failure. A QUIC version that violates the assumptions in this section therefore cannot be safely deployed with a load balancer that follows this specification. An updated or alternative version of this specification might address these shortcomings for such a QUIC version.

9. 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 3.1. 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. Due to the limitations of measures at QUIC layer, there is no scheme that strictly prevents linkability for all traffic patterns.

To see why, consider two limits. At one extreme, one client is connected to the server pool and migrates its address. An observer can easily link the two addresses, and there is no remedy at the QUIC layer.

At the other extreme, a very large number of clients are connected to each server, and they all migrate address constantly. At this limit, even an unencrypted server ID encoding is unlikely to definitively link two addresses.

Therefore, efforts to frustrate any analysis of server ID encoding have diminishing returns. Nevertheless, this specification seeks to minimize the probability two addresses can be linked.

9.1. 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.

9.2. Attackers between the load balancer and server

Attackers in this privileged position are intrinsically able to map two connection IDs to the same server. These algorithms ensure that two connection IDs for the same connection cannot be identified as such as long as the server chooses the first octet and any plaintext nonce correctly.

9.3. Multiple Configuration IDs

During the period in which there are multiple deployed configuration IDs (see Section 3.1), 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.

9.4. 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 4.3) 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 seven configurations distinguished by the config rotation bits (Section 3.1), exposing information about the target domain to the entire network; or
- * tenants can use the 0b111 codepoint 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 evaluate the privacy tradeoff by 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.

In the case that the administrating entity also controls a reverse proxy between the load balancer and the tenants, this entity generates the external CIDs, and there is no tradeoff.

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

9.5. 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.

9.6. 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 encrypted 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.

Servers are forbidden from generating linkable plaintext nonces, because 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.

9.7. Distinguishing Attacks

The Four Pass Encryption algorithm is structured as a 4-round Feistel network with non-bijective round function. As such, it does not offer a very high security level against distinguishing attacks, as explained in [Patarin2008]. Attackers can mount these attacks if they are in possession of $O(\text{SQRT}(\text{len}/2))$ pairs of ciphertext and known corresponding plain text, where "len" is the sum of the lengths of the Server ID and the Nonce.

The authors considered increasing the number of passes from 4 to 12, which would definitely block these attacks. However, this would require 12 round of AES decryption by load balancers accessing the CID, a cost deemed prohibitive in the planned deployments.

The attacks described in [Patarin2008] rely on known plain text. In a normal deployment, the plain text is only known by the server that generates the ID and by the load balancer that decrypts the content of the CID. Attackers would have to compensate by guesses about the allocation of server identifiers or the generation of nonces. These attacks are thus mitigated by making nonces hard to guess, as specified in Section 9.6, and by rules related to mixed deployments that use both clear text CID and encrypted CID, for example when transitioning from clear text to encryption. Such deployments MUST use different server ID allocations for the clear text and the encrypted versions.

These attacks cannot be mounted against the Single Pass Encryption algorithm.

9.8. Early deletion of load balancer connection state

Potential vulnerabilities related to heuristics that delete per-connection state are described in Section 6. Under certain assumptions about server configuration and fallback algorithm, this state might be critical to maintaining connectivity. Under other assumptions, the state provides robustness to improbable network events.

10. IANA Considerations

There are no IANA requirements.

11. References

11.1. Normative References

[NIST-AES-ECB]

Dworkin, M., "Recommendation for Block Cipher Modes of Operation: Methods and Techniques", NIST Special Publication 800-38A, 2021,
<<https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-38a.pdf>>.

[RFC8999] Thomson, M., "Version-Independent Properties of QUIC", RFC 8999, DOI 10.17487/RFC8999, May 2021,
<<https://www.rfc-editor.org/rfc/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/rfc/rfc9000>>.

11.2. Informative References

[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-25, 14 June 2025,
<<https://datatracker.ietf.org/doc/html/draft-ietf-tls-esni-25>>.

[Patarin2008]

Patarin, J., "Generic Attacks on Feistel Schemes - Extended Version", 2008,
<<https://eprint.iacr.org/2008/036.pdf>>.

[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/rfc/rfc2119>>.

[RFC4347] Rescorla, E. and N. Modadugu, "Datagram Transport Layer Security", RFC 4347, DOI 10.17487/RFC4347, April 2006,
<<https://www.rfc-editor.org/rfc/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/rfc/rfc6020>>.
- [RFC6066] Eastlake 3rd, D., "Transport Layer Security (TLS) Extensions: Extension Definitions", RFC 6066, DOI 10.17487/RFC6066, January 2011, <<https://www.rfc-editor.org/rfc/rfc6066>>.
- [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/rfc/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/rfc/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/rfc/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/rfc/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/rfc/rfc8484>>.
- [RFC9146] Rescorla, E., Ed., Tschofenig, H., Ed., Fossati, T., and A. Kraus, "Connection Identifier for DTLS 1.2", RFC 9146, DOI 10.17487/RFC9146, March 2022, <<https://www.rfc-editor.org/rfc/rfc9146>>.
- [RFC9147] Rescorla, E., Tschofenig, H., and N. Modadugu, "The Datagram Transport Layer Security (DTLS) Protocol Version 1.3", RFC 9147, DOI 10.17487/RFC9147, April 2022, <<https://www.rfc-editor.org/rfc/rfc9147>>.

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 "2023-07-14" {
  description
    "Updated to design in version 17 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..6";
    }
    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
        by nonce-length.";
    }

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

```

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.";
}

list server-id-mappings {
  key "server-id";
  description "Statically allocated Server IDs";

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

  leaf server-address {
    type inet:ip-address;
    mandatory true;
    description
      "Destination address corresponding to the server ID";
  }
}
}
}
}

```

A.1. Tree Diagram

This summary of the YANG models uses the notation in [RFC8340].

```

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 20a350d28b4203487d970b

```

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 0720b1d07b359d3c
1 ed793a51d49b8f5fab65 ee080dbf48
                                2fcc381bc74cb4fbad2823a3d1f8fed2
2 ed793a51d49b8f5f ee080dbf48c0d1e5
                                504dd2d05a7b0de9b2b9907afb5ecf8cc3
3 ed793a51d49b8f5fab ee080dbf48c0d1e55d
                                125779c9cc86beb3a3a4a3ca96fce4bfe0cdbc

```

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 [RFC9146] 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.

[RFC9146] 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 [RFC9147] 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, Greg Greenway, Jana Iyengar, Subodh Iyengar, Stefan Kolbl, Ladislav Lhotka, Jan Lindblad, Ling Tao Nju, Ilari Liusvaara, Kazuho Oku, Udip Pant, Zaheduzzaman Sarker, Ian Swett, Andy Sykes, Martin Thomson, Dmitri Tikhonov, Victor Vasiliev, Xingcan Lan, Yu Zhu, 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-20

- * Changed definition of Unroutable DCIDs, and rewrote sections on config failover and fallback routing to avoid misrouted connections.
- * Deleted text on dropping packets
- * Rewrote version invariance section

E.2. since draft-ietf-quic-load-balancers-19

- * Further guidance on multiple server processes/threads
- * Fixed error in encryption example.
- * Clarified fallback algorithms and known QUIC versions.

E.3. since draft-ietf-quic-load-balancers-18

- * Rearranged the output of the expand function to reduce CPU load of decrypt

E.4. since draft-ietf-quic-load-balancers-17

- * fixed regressions in draft-17 publication

E.5. since draft-ietf-quic-load-balancers-16

- * added a config ID bit (now there are 3).

E.6. since draft-ietf-quic-load-balancers-15

- * aasvg fixes.

E.7. since draft-ietf-quic-load-balancers-14

- * Revised process demultiplexing text
- * Restored lost text in Security Considerations
- * Editorial comments from Martin Thomson.

- * Tweaked 4-pass algorithm to avoid accidental plaintext similarities
- E.8. since draft-ietf-quic-load-balancers-13
- * Incorporated Connection ID length in argument of truncate function
 - * Added requirements for codepoint 0b11.
 - * Describe Distinguishing Attack in Security Considerations.
 - * Added non-normative language about server process demultiplexers
- E.9. since draft-ietf-quic-load-balancers-12
- * Separated Retry Service design into a separate draft
- E.10. since draft-ietf-quic-load-balancers-11
- * Fixed mistakes in test vectors
- E.11. 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.12. 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.13. since draft-ietf-quic-load-balancers-08
- * Eliminate Dynamic SID allocation
 - * Eliminated server use bytes

E.14. 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.15. 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.16. 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.17. 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.18. since-draft-ietf-quic-load-balancers-03

- * Improved Config Rotation text
- * Added stream cipher test vectors
- * Deleted the Obfuscated CID algorithm

E.19. 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.20. 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.21. since-draft-ietf-quic-load-balancers-00

- * Removed in-band protocol from the document

E.22. Since draft-duke-quic-load-balancers-06

- * Switch to IETF WG draft.

E.23. 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.24. Since draft-duke-quic-load-balancers-04

- * Added standard for retry services

E.25. 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.26. 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.27. Since draft-duke-quic-load-balancers-01

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

E.28. 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