Workgroup: IDR Workgroup
Internet-Draft: draft-retana-idr-bgp-quic-01
Published: 12 March 2023
Intended Status: Standards Track
Expires: 13 September 2023
Authors: A. Retana
Futurewei Technologies
Y. Qu
Futurewei Technologies
J. Haas
Juniper Networks
S. Chen
Huawei Technologies
J. Tantsura
Microsoft

BGP over QUIC

Abstract

This document defines the use of QUIC as BGP transport protocol.

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 13 September 2023.

Copyright Notice

Table of Contents

1. Introduction

The Border Gateway Protocol (BGP) [RFC4271] is the routing protocol used to exchange routing and reachability information among autonomous systems, and it uses TCP as its transport protocol to provide reliable packet communication. BGP establishes peer relationships between routers using a TCP session on port 179.

The Multiprotocol Extensions for BGP-4 (MP-BGP) [RFC4760] allow BGP to carry information for multiple Network Layer protocols. However, only a single TCP connection can reach the Established state between a pair of peers [RFC4271]. As a consequence, an error related to a

particular Network Layer protocol may result in the termination of the connection for all.

QUIC [RFC9000] is a UDP-based multiplexed and secure transport protocol that provides connection-oriented and stateful interaction between a client and server. It can provide low latency and encrypted transport with resilient connections.

This document specifies the procedures for BGP to use QUIC as a transport protocol (Section 3), including error handling (Section 4). Changes to the BGP Finite State Machine (FSM) [RFC4271] are described in Section 5.

1.1. Requirements Language

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

2. Terminology

This document relies on the terminology used in [RFC4271] and [RFC9000]. Familiarity with those documents is expected.

3. BGP over QUIC (BoQ)

BGP over QUIC (BoQ) replaces only the transport layer of BGP. The BGP protocol specification remains backward compatible.

Before two BGP speakers start exchanging routing information, they need to establish a BGP session. It is established in two phases:

1. Establish a transport layer connection. Section 3.1 specifies the BoQ connecition establishment.

2. Establish a BGP session. After a transport-layer connection is established, BGP peers exchange protocol messages as specified in [RFC4271]. Section 3.2 specifies a mechanism that allows multiple BGP sessions in a single QUIC connection.

[RFC4271] defines the BGP FSM operation, including transport connection conflict detection and resolution. BoQ follows the same definitions except where explicit modifications are defined in this document.

3.1. BoQ Connection Establishment

BoQ uses QUIC version 1 as the underlying transport. Other QUIC versions that meet the definition of a compatible version [I-D.ietf-quic-version-negotiation] with version 1 MAY be used. The use of incompatible QUIC versions, as defined in [I-D.ietf-quic-version-negotiation], may be specified in future documents.

QUIC connections are established as described in [RFC9000]. During connection establishment, a BGP speaker SHOULD use UDP port 179 and MUST select the Application-Layer Protocol Negotiation (ALPN) [RFC7301] token "bgp4" in the TLS handshake. Support for other application-layer protocols MUST NOT be offered in the same handshake. A connection MUST be closed if the ALPN token is not as indicated, or if other application-layer protocols are offered in the same handshake.

After a QUIC connection is established, the first message sent by each endpoint is a BGP OPEN message, which can be used to indicate whether the speaker is capable of supporting more than one QUIC stream to exchange BGP messages (Section 3.2). The message format and framing is unchanged [RFC4271].

3.2. Multiple BGP Sessions

In QUIC, application protocols exchange information using streams. Each stream is a separate unidirectional or bidirectional channel of "order stream of bytes." Moreover, each stream has flow control which limits bytes sent on a stream, together with flow control of the connection. Multiple streams can be multiplexed onto an underlying connection.

3.2.1. Multiple BGP Sessions Using QUIC Streams

This document specifies a mechanism to establish multiple BGP sessions using QUIC streams, one session per stream. An implementation can assign one or more Network Layer protocols to a BGP session.

A QUIC stream is created by sending a BGP OPEN message, and each stream MUST be bidirectional as described in Section 2.1 of [RFC9000]. In addition, the corresponding stream MUST end (clean termination) as described in Section 2.4 of [RFC9000] when a BGP session is terminated.

Section 3.2.6 describes the Connection Collision Detection procedure to be used with streams. Each BGP session operates independently, which means critical conditions (such as a malformed message) in one session won't affect others.

3.2.2. MultiStream Capability

The MultiStream Capability (MSC) is defined to indicate that a BGP speaker supports multiple sessions as specified in this document. The capability [RFC5492] is defined as follows:

Capability code (1 octet): TBD1

Capability length (1 octet): 1

Capability value (1 octet): flag field reserved.

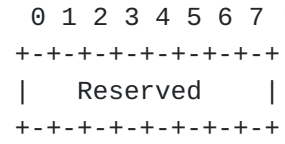

Flags: bitfield - MUST be set to zero and ignored by the receiver.

The MSC only applies when using BoQ. It MUST be included in all OPEN messages. It MUST be ignored otherwise.

BGP multiple session support defined in Section 3.2 applies only if both peers advertise the MSC during the establishment of the "initial session." In particular, if a peer that advertises the MSC doesn't receive an OPEN message with the MSC from its peer, it SHOULD NOT terminate the session.

Using the MSC allows peers to establish multiple BGP sessions, one per QUIC stream. Each new BGP session is established using a separate OPEN message [RFC4271] and MUST include the MSC. If both peers exchange the MSC in the "initial session," they MUST include it when establishing other sessions. Otherwise, the new session MUST be terminated, and the Error Subcode MUST be set to MultiStream Conflict (TBD2), defined in Section 4.

Once a BGP session is established, it follows the procedures specified in [RFC4271].

3.2.3. The Control Stream and Function Streams

If both peers support MSC, the "initial session" creates a control stream for the BGP connection. The OPEN messages exchanged contrain information such as its AS number and the BGP protocol version it supports, if the peers agree on the parameters for the session, they will begin to send BGP KEEPALIVE messages to each other.

If the entire BGP connection needs to be reset for any reason, such as a configuration change or a network outage, a notification is

sent over the control stream to inform the other router that the connection is being reset.

If for whatever reason the control stream is closed, the QUIC connection needs to be terminated using a CONNECTION_CLOSE frame, and an error message (TBD) should be included to indicate that the connection has been terminated by BGP. If there are other open streas, they are implicitly clsed when the connection is closed.

In addition to the control stream, there may be other function streams that are used to exchange specific types of routing information. For example, one AFI/SAFI may be mapped to one function stream. These function streams cannot be established until the control stream has reached an established state.

Each function stream has its own keepalive messages with its own timers. These timers are used to ensure that the session stays alive even if no routing updates are being exchanged. Non-routing related function streams, such as BGP FLowSpec, may have longer keepalive timers, for example 120 seconds, as they do not need to exchange routing updates as frequently as other function streams and can run in a relatively lower priority.

A single QUIC stream provides ordered and reliable delivery, however there is no guarantee of transmission and deliver order across streams. Therefore, if specific data from one stream needs to be received before data from other streams, this requirement must be accomplished through BGP..

3.2.4. Stream Priorites

As defined in [RFC9000], a QUIC implementation SHOULD provide ways in which an application can indicate the relative priority of streams.

For a BGP implementation utilizing QUIC as its transport protocol with MultiStream Capability, it MUST support a prioritization mechanism for BGP streams. This is essential for ensuring that critical routing information can be transmitted with higher priority compared to non-routing information.

How to implement the supported priorities using QUIC congestion control at connection level, stream level flow control, and packetization are out of the scope of this document.

3.2.5. Modifications to the BGP FSM

The modifications to the BGP FSM are described in Section 5. For simplicity and security reason, it is suggested that 1-RTT is used.

BGP multi-session support doesn't modify the BGP FSM, but the collision handling procedure should be replaced with the procedure described below.

3.2.6. BGP Session Establishment and Collision Avoidance

Before creating a new session, a BGP speaker should check that no session exists for the same Network Layer protocol(s). If a session already exists, the BGP speaker SHOULD NOT attempt to create a new one.

If a pair of BGP speakers try to establish a BGP session with each other simultaneously, then two parallel sessions will be formed. In the case of BoQ, the IP addresses of the connection cannot be used to resolve collisions when using multiple streams.

To avoid connection collisions, a session is identified by the My Autonomous System and BGP Identifier fields pair in the OPEN message. In this context, a connection collision is the attempt to open a BGP session for which the set of Network Layer protocols is the same. One of the connections MUST be closed.

The connection collision is resolved using the extension specified in [RFC6286]. In other words, the session with the higher-valued BGP Identifier is preserved [RFC4271]. If the BGP Identifiers are identical, then the session with the larger ASN is preserved [RFC6286].

Upon receiving an OPEN message, the local system MUST examine all of its sessions in the OpenConfirm state. A BGP speaker MAY also examine sessions in an OpenSent state if it knows the BGP Identifier of the peer by means outside of the protocol. If among these sessions, there is one to a remote BGP speaker whose BGP Identifier and ASN pair equals the one in the OPEN message, and this session collides with the connection over which the OPEN message is received, then the local system performs the following collision resolution procedure:

1) The BGP Identifier of the local system is compared to the BGP Identifier of the remote system (as specified in the OPEN message). Comparing BGP Identifiers is done by converting them to host byte order and treating them as 4-octet unsigned integers.

2) If the value of the local BGP Identifier is less than the remote one, the local system closes the BGP connection that already exists (the one that is already in the OpenConfirm state) and accepts the BGP connection initiated by the remote system.

2a) Otherwise, the local system closes the newly created BGP connection (the one associated with the recently received OPEN

message) and continues to use the existing one (the one that is already in the OpenConfirm state).

3) If the BGP Identifiers of the peers involved in the connection collision are identical, then the session initiated by the BGP speaker with the larger AS number is preserved.

Unless allowed via configuration, a connection collision with an existing BGP session in the Established state causes the closing of the newly created session.

Closing the BGP session (that results from the collision resolution procedure) is accomplished by sending the NOTIFICATION message with the Error Code Cease, Subcode Connection Collision Resolution (7) [RFC4486].

The remainder of the process is as specified in [RFC4271].

4. Error Handling

BoQ error handling involves the following three types of errors:

(1) QUIC error: Includes stream error and connection error [RFC9001]. In some cases, a stream error may cause a connection error. For example, if an operation error occurs on all streams, the connection error should be triggered to close the connection.

(2) TLS alert: In [RFC9001],a QUIC endpoint MUST treat any alert from TLS as if it were at the "fatal" level. For TLS alerts, this includes replacing any alert with a generic alert, such as handshake_failure (0x128 in QUIC).

(3) BGP error: If an error occurs in BGP processing [RFC4271], it can be mapped to the following BoQ Error Codes[RFC9000].

This document defines some of the following BoQ Error Codes:

(1) BOQ_NO_ERROR (0x00): No error. This is used when the connection or stream needs to be closed, but there is no error to signal.

(2) BOQ_INTERNAL_ERROR (0x01): The BoQ implementation encountered an internal error and is incapable of continuing the stream or the connection.

4.1. Error Handling with MultiSteam Support

OPEN message error handling is defined in section 6.2 of [RFC4271]. This document introduces the following OPEN Message Error subcodes:

TBD2 - MultiSession Conflict - Used if the MSC is exchanged by both peers in the "initial session" but is not present when establishing a new session.

TBD3 - Session Capability Mismatch - Used if a BGP speaker terminates a session in the case where it sends an OPEN message with the MSC but receives an OPEN message without it.

TBD4 - Network Layer Protocol Mismatch - Used if a BGP session has already been established for a signaled Network Layer Protocol, either individually or as part of a set.

Section 3.2.2 recommends not terminating a session when only one peer supports the MSC. If such a BGP speaker does terminate the session, the Error Subcode MUST be set to Session Capability Mismatch (TBD3).

Any individual BGP session can be terminated as specified in [RFC4486]. If multiple sessions are to be terminated, then the procedure MUST be followed for each one.

4.2. Session closure

QUIC provides three ways to close a connection(see [RFC9000] Section 10):

(1) Idle timeout

(2) Immediate Close

(3) Stateless Reset

When the idle timer expires, the connection is closed immediately. Idle timeout can be calculated using the following formula:

idle_timeout=MAX(min_idle_timeout, 3*PTO)

The PTO is a time that the sender should wait for an acknowledgment of a sent packet. For a calculation method, refer to [RFC9002] Section 6.2.1.

When establishing a QUIC connection, the transmission parameter max_idle_timeout is used. Endpoints advertise local idle_timeout to each other. If no max_idle_timeout advertisement is received from the remote end, the remote idle_timeout is set to a value of 0.

Based on the values of local idle_timeout and remote idle_timeout, there are three possible scenarios:

(1) If both the values are 0, disable the idle timeout function.

(2) If there is only one value 0, set min_idle_timeout to a non-zero value in between.

(3) If neither value is 0, set min_idle_timeout to the smaller value.

Two options are available for the idle timer during BGP session establishment. Option 1 is recommended by default.

Option 1: Set this parameter to 0, indicating that idle timeout is disabled.

Option 2: The value must be greater than the value of BGP HoldTimer. It is recommended that the value be greater than five times the value of BGP HoldTimer.

5. BGP Finite State Machine

5.1. Optional Session Attributes

This document adds two optional Session attributes to the list in Section 8 of [RFC4271]:

14) PassiveQUICEstablishment
15) TrackQUICState

Section 8.1.1 of [RFC4271] describes the linkage between the FSM functionality, events, and optional session attributes. When using BoQ, Group 3 (TCP processing) is replaced with:

Group 3: QUIC processing

Optional Session Attributes: PassiveQUICEstablishment, TrackQUICState

Option 1: PassiveQUICEstablishment

Description: This option indicates that the BGP FSM will passively wait for the remote BGP peer to establish the BGP QUIC connection. The local node is a QUIC server [RFC9000].

Value: TRUE or FALSE

Option 2: TrackQUICState

Description: The BGP FSM tracks the end result of a QUIC connection attempt rather than individual QUIC messages. Optionally, the BGP FSM can support additional interaction with the TCP connection negotiation.

Value: TRUE or FALSE

5.2. FSM Event

QUIC directly encapsulates the handshake process of TLS 1.3 [RFC8446]. In addition, QUIC requires that all packets must be explicitly acknowledged. Therefore, QUIC defines the end state of two connection establishment [RFC9001]

(1) Handshake Complete: TLS 1.3 has successfully completed the handshake.

(2) Handshake Confirmed: The QUIC has successfully completed the handshake.

On the QUIC client, the state is Handshake Complete and then Handshake Confirmed. On the QUIC server, the two states are reached at the same time.

The transport layer events for BoQ FSM are defined as follows :

Event 29: ManualStart_with_PassiveQuicEstablishment

Definition: Local system administrator manually starts the peer connection, but has PassiveQuicEstablishment enabled.

Status: Optional, depending on local system

Optional Attribute Status:

1) The PassiveTcpEstablishment attribute SHOULD be set to TRUE if this event occurs.

2) The DampPeerOscillations attribute SHOULD be set to FALSE when this event occurs.

Corresponding TCP events: Event 4

Event 30: AutomaticStart_with_PassiveQuicEstablishment

Definition: Local system automatically starts the BGP connection with the PassiveQuicEstablishment enabled.

Status: Optional, depending on local system

Optional Attribute Status:

1) The AllowAutomaticStart attribute SHOULD be set to TRUE.

2) The PassiveTcpEstablishment attribute SHOULD be set to TRUE.

3) If the DampPeerOscillations attribute is supported, the DampPeerOscillations SHOULD be set to FALSE.

Corresponding TCP events: Event 5

Event 31: AutomaticStart_with_DampPeerOscillations_and_PassiveQuicEstablishment

Definition: Local system automatically starts the BGP peer connection with peer oscillation damping enabled and PassiveQuicEstablishment enabled. The exact method of damping persistent peer oscillations is determined by the implementation and is outside the scope of this document.

Status: Optional, depending on local system

Optional Attribute Status:

1) The AllowAutomaticStart attribute SHOULD be set to TRUE.

2) The DampPeerOscillations attribute SHOULD be set to TRUE.

3) The PassiveTcpEstablishment attribute SHOULD be set to FALSE.

Corresponding TCP events: Event 7

Event 32: QuicConnection_Valid

Definition: This parameter is applicable only to the QUIC server. It indicates that the Handshake Confirmed state is reached.

Status: Optional

Optional Attribute Status: 1) The TrackTcpState attribute SHOULD be set to TRUE if this event occurs.

Corresponding TCP events: Event 14

Event 33: Quic_CR_Invalid

Definition: This parameter applies only to the QUIC server and indicates that an invalid QUIC connection request is received. Initial packets with invalid source addresses or port

numbers,invalid destination addresses or port numbers or version negotiation or address validation fails.

Status: Optional

Optional Attribute Status: 1) The TrackTcpState attribute should be set to TRUE if this event occurs.

Corresponding TCP events: Event 15

Event 34: Quic_CR_Acked

Definition: This parameter applies only to the QUIC client. It indicates that an Initial ACK message is received from the QUIC server and an Initial/Handshake message is sent to the QUIC server. Note: When this event is received, the QUIC client has reached the Handshake Complete state.

Status: Mandatory

Corresponding TCP events: Event 16

Event 35: QuicConnectionConfirmed

Definition: This parameter applies to both QUIC client and QUIC server, indicating that the Handshake Confirmed state has been reached.

Status: Mandatory

Corresponding TCP events: Event 17

Event 36: QuicConnectionFails

Definition: This parameter applies to both the QUIC client and the QUIC server. It indicates that an error occurs in the QUIC handshake before the system enters the Handshake Confirmed state.

Status: Mandatory

Corresponding TCP events: Event 18

6. Operational Considerations

6.1. Using BoQ

The decision to use BoQ instead of the TCP-based mechanism defined in [RFC4271] is an operational decision and out of the scope of this document. An implementation MUST provide a configuration mechanism to enable BoQ on a per-peer basis. More granularity (per Network

Layer protocol, for example) is not recommended as it may increase the operational complexity.

Connectivity problems (e.g., blocking UDP) can result in a failure to establish a QUIC connection; BGP speakers SHOULD attempt to establish a TCP-based BGP session in this case.

6.2. BGP Multi Session Backward Compatibility

A BGP speaker that doesn't understand the MSC will ignore it [RFC5492]. Section 3.2.2 recommends not terminating a session when only one peer supports the MSC.

6.3. BGP Multi Session Prioritization

One of the drawbacks of a single BGP session is that control plane messages for all supported Network Layer protocols use the same connection, which may cause resource contention.

QUIC [RFC9000] does not provide a mechanism for exchanging prioritization information. Instead, it recommends that implementations provide ways for an application to indicate the relative priority of streams, in this case, mapped to BGP sessions. An operator should prioritize BGP sessions (streams) that carry critical control plane information if the functionality is available. The definition of this functionality and the determination of the importance of a BGP session are both outside the scope of this document.

An example implementation is to have four priority (0-3) defined, and smaller number means higher priority. Each AFI/SAFI should be assigned a default priority and optional configuration to modify the default value. For example, IPv4 and IPv6 unicast AFI/SAFI (1/1 and 2/1) may have priority of 1, while BGP-LS (16388/71 and 16388/72) may have a priority of 3, and BGP FlowSpec (1/133 and 1/134) may have a priority of 4.

6.4. Configurations

For BGP multi session, a configuration command SHOULD be implemented to allow grouping of some AFI/SAFIs into one session.

7. Security Considerations

This document replaces the transport protocol layer of BGP from TCP to QUIC. It does not modify the basic protocol specifications of BGP, and therefore does not introduce new security risks to the basic BGP protocol. The non-TCP-related considerations of [RFC4271], [RFC4272], and [RFC7454] apply to the specification in this document.

BoQ enhances transport-layer security for BGP sessions, refer to [RFC7454] :

(1) Supports QUIC server identity authentication.

(2) (Optional) Supports QUIC client identity authentication.

(3) Confidentiality protection of BGP messages is supported. All BGP messages are encrypted for transmission.

(4) Supports integrity protection for BGP messages.

The use of a specific UDP port number and an ALPN token Section 3.1 protects a BGP Speaker from attempts to establish an unexpected BGP session. Additionally, all packets directed to UDP port 179 on the local device and sourced from an address not known or permitted to become a BGP neighbor SHOULD be discarded.

With BGP multi session support using QUIC streams, it separates the control plane traffic over multiple sessions, the effect of a session-based vulnerability is reduced; only a single session is affected and not the whole connection. The result is increased resiliency.

On the other hand, a high number of BGP sessions may result in higher resource utilization and the risk of depletion. Also, more sessions may imply additional configuration and operational complexity. However, this risk is mitigated by the fact that BGP sessions typically require explicit configuration by the operator.

8. IANA Considerations

8.1. UDP Port for BoQ

IANA is requested to add a reference to [this document] for the UDP port 179 entry in the "Service Name and Transport Protocol Port Number Registry".

8.2. Registration of the BGP4 Identification String

This document creates a new registration for the identification of BGP [RFC4271] in the "TLS Application-Layer Protocol Negotiation (ALPN) Protocol IDs" registry.

The "bgp4" string identifies BGP-4 [RFC4271]:

Protocol: BGP-4

Identification Sequence: 0x62 0x67 0x70 0x34 ("bgp4")

Specification: This document

8.3. Multiple Streams

IANA is asked to assign a new Capability Code for the MultiStream Capability (Section 3.2.2) as follows:

Value	Description	Reference	Change Controller
TBD1	MultiStream Capability	[This Document]	IETF

Table 1: MultiStream Capability

8.4. Error Codes

IANA is asked to assign three values from the OPEN Message Error subcodes registry as follows:

Value	Name	Reference
TBD2	MultiSession Conflict	[This Document]
TBD3	Session Capability Mismatch	[This Document]
TBD4	Network Layer Protocol Mismatch	[This Document]

Table 2

IANA is asked to assign two values from the Cease NOTIFICATION Message Error subcodes registry as follows:

Value	Name	Reference
BOQ_NO_ERROR	Stream Closed No Error	[This Document]
BOQ_INTERNAL_ERROR	Stream Internal Error	[This Document]

Table 3

9. Acknowledgement

This document merges and replaces [I-D.chen-idr-bgp-over-quic] and [I-D.retana-idr-bgp-quic-stream]. The authors acknowledge the contributions made by the authors and contributors of those documents.

This document references the text and procedures defined in [I-D.ietf-idr-bgp-multisession], and we are grateful for their contributions.

The authors would like to thank xx for review and comments.

10. References

10.1. Normative References

[I-D.ietf-quic-version-negotiation] Schinazi, D. and E. Rescorla, "Compatible Version Negotiation for QUIC", Work in

Progress, Internet-Draft, draft-ietf-quic-version-negotiation-11, 11 October 2022, <https://datatracker.ietf.org/doc/html/draft-ietf-quic-version-negotiation-11>.

[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>.

[RFC4271] Rekhter, Y., Ed., Li, T., Ed., and S. Hares, Ed., "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, DOI 10.17487/RFC4271, January 2006, <https://www.rfc-editor.org/info/rfc4271>.

[RFC4486] Chen, E. and V. Gillet, "Subcodes for BGP Cease Notification Message", RFC 4486, DOI 10.17487/RFC4486, April 2006, <https://www.rfc-editor.org/info/rfc4486>.

[RFC5492] Scudder, J. and R. Chandra, "Capabilities Advertisement with BGP-4", RFC 5492, DOI 10.17487/RFC5492, February 2009, <https://www.rfc-editor.org/info/rfc5492>.

[RFC6286] Chen, E. and J. Yuan, "Autonomous-System-Wide Unique BGP Identifier for BGP-4", RFC 6286, DOI 10.17487/RFC6286, June 2011, <https://www.rfc-editor.org/info/rfc6286>.

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

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

[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.chen-idr-bgp-over-quic] Chen, S., Zhang, Y., Wang, H., and Z. Li, "BGP Over QUIC", Work in Progress, Internet-Draft, draft-chen-idr-bgp-over-quic-00, 3 June 2021, <https://datatracker.ietf.org/doc/html/draft-chen-idr-bgp-over-quic-00>.

[I-D.ietf-idr-bgp-multisession] Scudder, J., Appanna, C., and I. Varlashkin, "Multisession BGP", Work in Progress, Internet-Draft, draft-ietf-idr-bgp-multisession-07, 13 September 2012, <https://datatracker.ietf.org/doc/html/draft-ietf-idr-bgp-multisession-07>.

[I-D.retana-idr-bgp-quic-stream] Retana, A., Qu, Y., and J. Tantsura, "Use of Streams in BGP over QUIC", Work in Progress, Internet-Draft, draft-retana-idr-bgp-quic-stream-02, 11 May 2022, <https://datatracker.ietf.org/doc/html/draft-retana-idr-bgp-quic-stream-02>.

[RFC4272] Murphy, S., "BGP Security Vulnerabilities Analysis", RFC 4272, DOI 10.17487/RFC4272, January 2006, <https://www.rfc-editor.org/info/rfc4272>.

[RFC4760] Bates, T., Chandra, R., Katz, D., and Y. Rekhter, "Multiprotocol Extensions for BGP-4", RFC 4760, DOI 10.17487/RFC4760, January 2007, <https://www.rfc-editor.org/info/rfc4760>.

[RFC7454] Durand, J., Pepelnjak, I., and G. Doering, "BGP Operations and Security", BCP 194, RFC 7454, DOI 10.17487/RFC7454, February 2015, <https://www.rfc-editor.org/info/rfc7454>.

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

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

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

Authors' Addresses

Alvaro Retana
Futurewei Technologies
2330 Central Expressway
Santa Clara, CA 95050
United States of America

Email: aretana@futurewei.com

Yingzhen Qu

Futurewei Technologies
2330 Central Expressway
Santa Clara, CA 95050
United States of America

Email: yingzhen.qu@futurewei.com

Jeffrey Haas
Juniper Networks

Email: jhaas@pfrc.org

Shuanglong Chen
Huawei Technologies
No.156 Beiqing Rd.
Beijing
100095
China

Email: chenshuanglong@huawei.com

Jeff Tantsura
Microsoft
United States of America

Email: jefftant.ietf@gmail.com