



# Use of Streams in BGP over QUIC

**draft-retana-idr-bgp-quic-stream**

---

IETF 112 November 2021

Alvaro Retana  
Yingzhen Qu  
Jeff Tantsura

# Why BGP multisession?

---

A single TCP connection between a pair of peers [RFC4271].

As pointed out by draft-ietf-bgp-multisession:

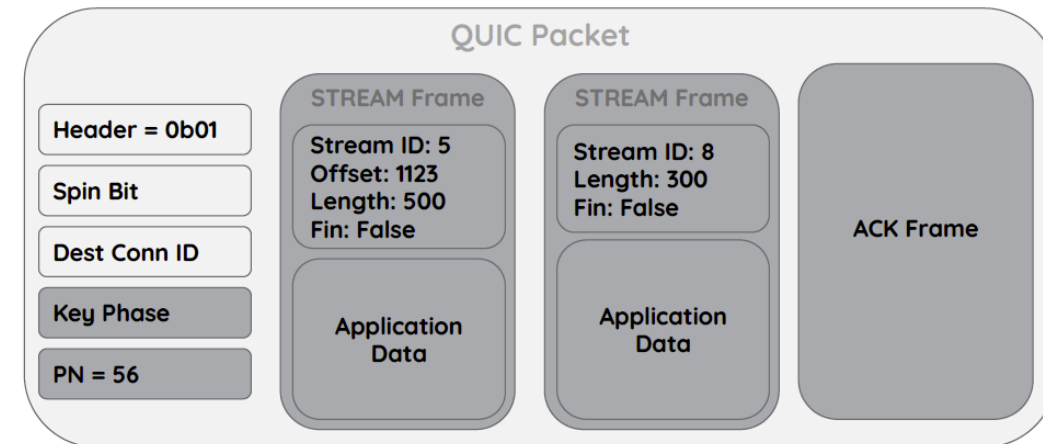

- A common criticism of BGP is the fact that most malformed messages cause the session to be terminated. While this behavior is necessary for protocol correctness, one may observe that the protocol machinery of a given implementation may only be defective with respect to a given AFI/SAFI. Thus, it would be desirable to allow the session related to that family to be terminated while leaving other AFI/SAFI unaffected.
- A second criticism of BGP is that it is difficult or in some cases impossible to manage control plane resource contention when BGP is used to support diverse services over a single session. In contrast, if a single BGP session carries only information for a single service (or related set of services) it may be easier to manage such contention.

# QUIC Streams

QUIC [RFC9000] is a UDP-based multiplexed and secure transport protocol.

In QUIC, application protocols exchange information via streams, and multiple streams can be multiplexed on to an underlying connection:

- Each stream is a separate unidirectional or bidirectional channel
- Each stream has flow control, as well as flow control of the connection.

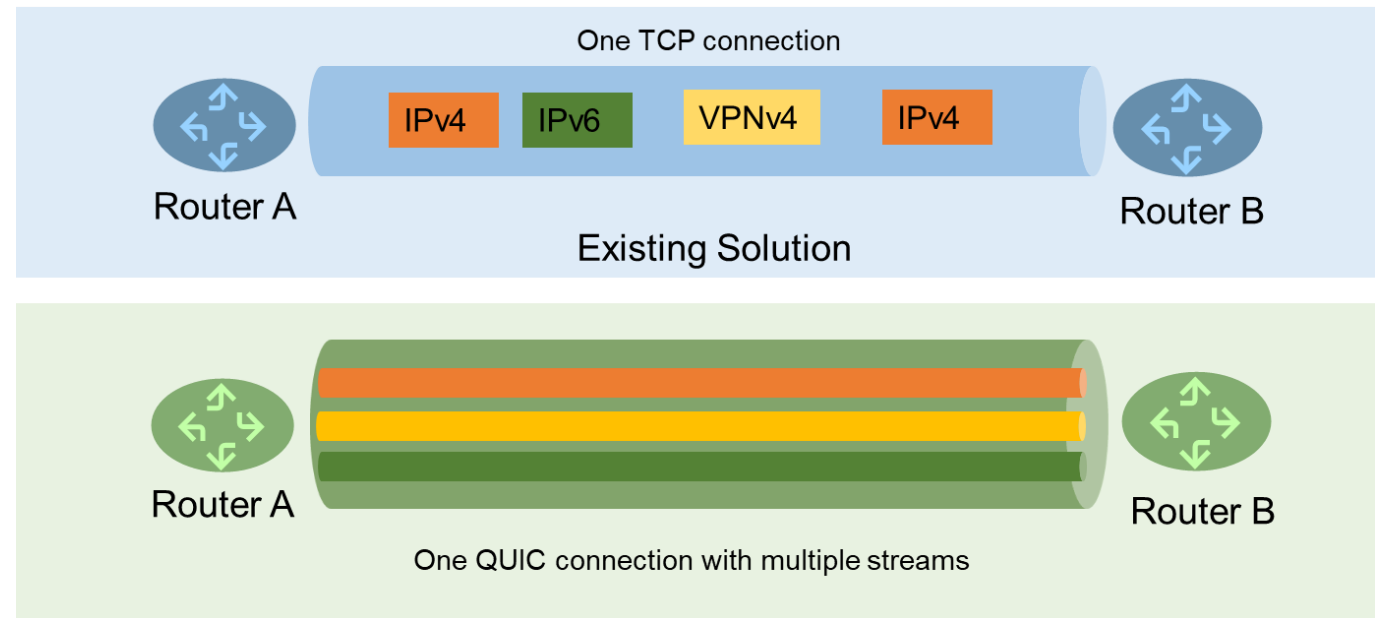


# BGP Multisession using QUIC Streams

Draft-chen-idr-bgp-over-quic defines the mechanism of BGP over QUIC and different options to map streams.

This proposal specifies a complementary and backwards compatible mechanism to establish multiple BGP sessions using QUIC streams.

- Each stream is created by sending a BGP OPEN message.
- Each stream MUST be bidirectional.
- Each BGP session operates independently, which means critical conditions in one session won't affect others.



# BGP MultiStream Capability

---

The MultiStream Capability (MSC) is defined to indicate that a BGP speaker supports multisession.

Using the MSC allows peers to establish multiple BGP sessions, one per QUIC stream.

Once a BGP session is established, it follows the procedures specified in RFC 4271.

The capability [[RFC5492](#)] is defined as follows:

Capability code (1 octet): TBD1

Capability length (1 octet): 1

Capability value (1 octet): flag field reserved.

```
 0 1 2 3 4 5 6 7
+---+---+---+---+
|   Reserved   |
+---+---+---+---+
```

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

# Error Handling

---

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

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.

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.

# Collision Avoidance

---

Check existing sessions before creating a new session

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].

Please refer to the draft for details.

# Next Steps

---

Continue to refine

Review and comments are welcome!

Collaborators are welcome!





**Thank you**

---

