

BGP over QUIC

[draft-retana-idr-bgp-quic - BGP over QUIC](#)

Alvaro Retana, Yingzhen Qu, **Jeff Haas**

Shuanglong Chen, Jeff Tantsura

Background - QUIC

- QUIC is a new UDP-based multiplexed and secure transport protocol defined in RFC 9000.

<https://datatracker.ietf.org/wg/quic/about/>

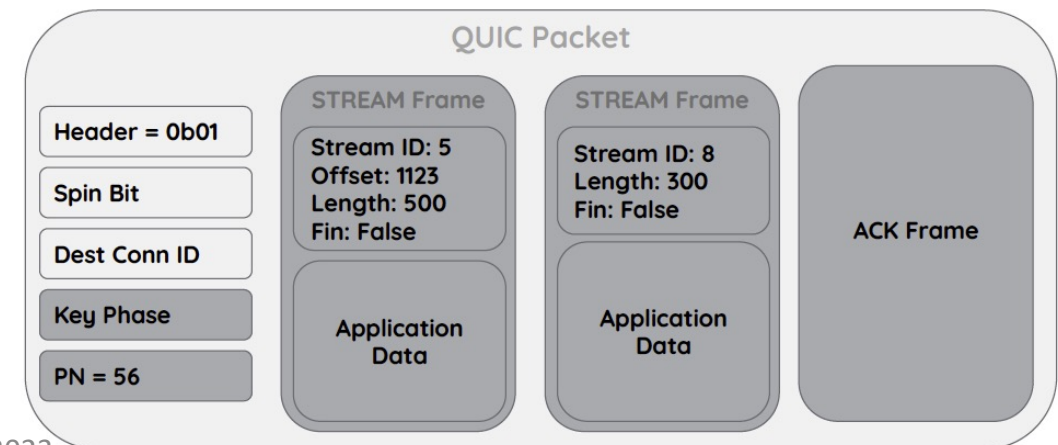

Low latency

Encrypted transport

Resilient connections (18-byte connection IDs)

- Multiple streams can be multiplexed onto an underlying connection

Each stream is a separate “ordered stream of bytes.
Stream flow control limits bytes sent on a stream.



Why BGP over QUIC

- High Resiliency

With multi-channel channel BGP support:

each channel can be reset independently without impacting other channels. draft-ietf-idr-bgp-multisession defines BGP over multiple TCP sessions but achieved very limited implementation and deployment.

Allowing each channel to transmit data concurrently. This helps avoid head-of-line blocking and improve overall performance.

- Enhanced Security

QUIC provides transport security with build-in encryption and authentication. This helps protect against spoofing attacks etc.

- Connection Migration

QUIC supports seamless migration of connections between different network interfaces or IP addresses.

Terminologies

- BoQ, Multi-channel BGP using QUIC: Running the BGP protocol over multiple QUIC streams as defined in this document.
- QUIC connection: A transport-layer connection between two endpoints using QUIC [RFC9000].
- QUIC streams: A bidirectional or unidirectional byte stream provided by the QUIC transport [RFC9000].
- BGP channel: Instance of BGP protocol state machine mapped to specific QUIC stream.
- BGP control channel: a channel dedicated to transmitting the session control data, which is implemented as a bidirectional stream.
- BGP function channel: BGP per AFI/SAFI channel, which is implemented asymmetrically as unidirectional streams.

Establish BoQ Connection

- Establish a transport layer connection. TLS 1.3 is integrated with QUIC. The TLS authentication parameters used for this connection are out of scope.

A BoQ speaker **MUST** select the Application-Layer Protocol Negotiation (ALPN) [RFC7301] token "boq" in the TLS handshake. Support for other application-layer protocols **MUST NOT** be offered in the same handshake.

- Establish a BoQ session over this transport connection. Defined in the draft.

Control Channel

- In one BoQ session, there is one control channel and multiple function channels.
- The control channel is a client-initiated ***bidirectional QUIC stream*** with stream ID 0. The capabilities carried in the OPEN message for the control channel are the BoQ connection specific parameter, such as BGP Role Capability.

Function Channels

- Per-AFI/SAFI function channels are used to exchange routing information. They are ***unidirectional streams*** and can be client or server initiated. Only one function channel per AFI/SAFI exists from one BoQ speaker to another.
- Function channels are ***asymmetrically*** created after the control channel reaches established state. A BoQ peer, each can create a function channel to implement symmetric route advertisements.
- Each function channel is created independently and can be ***reset in case of error without impacting other channels.***

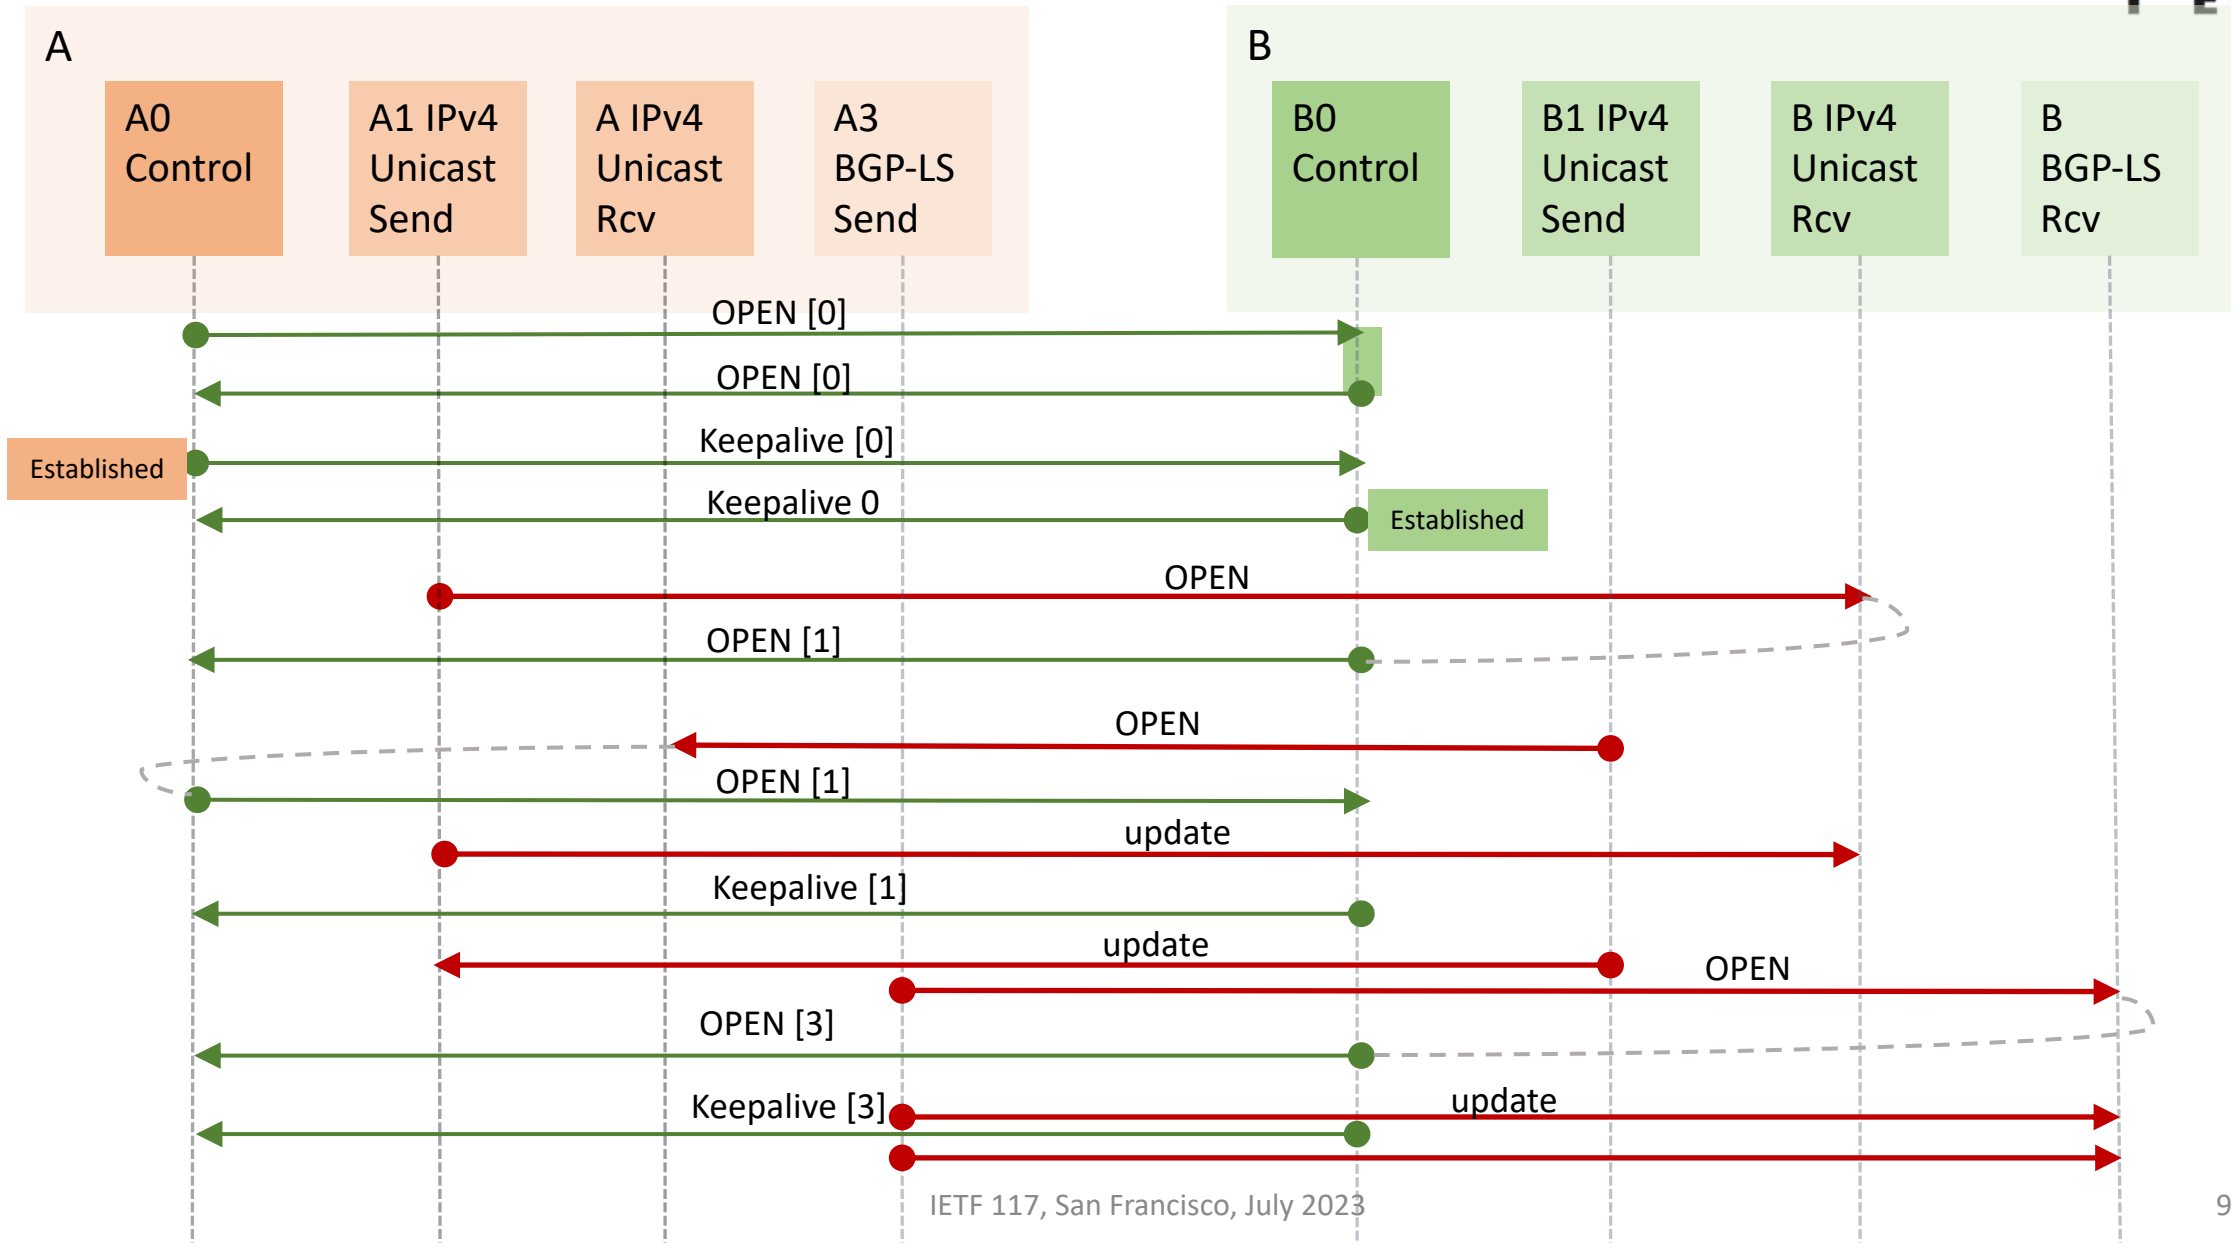
BoQ Framing Layer

- BoQ messages are carried by QUIC STREAM frames.
- There are BoQ messages sent in the control channel that are meant for a function channel and need to have the corresponding function channel/stream ID.
- The following frame types are defined with the frame payload carries the BGP message.

```
BoQ Data Frame {  
    Type (16) = 0,  
    Length (16),  
    Frame Payload (..)  
}
```

```
BoQ Control Data Frame {  
    Type (16) = 1,  
    Length (16),  
    Stream ID (62),  
    Padding (2) = 0,  
    Frame Payload (..)  
}
```


Multi-Channel BGP over QUIC



Channel Reset and Coordination

- NOTIFICATIONS are always sent over the control channel.
- Hold timer in the control channel SHOULD be smaller than those in function channels.
- QUIC doesn't guarantee transmission and delivery orders across streams. AFI/SAFI coordination must be accomplished through BGP.
Function channels for critical routing information should be given higher priority compared to non-routing information.

BGP over QUIC Capability

- QUIC supports connection migration, a feature considered for BoQ in the future. However, only the client side can move.
- An explicit configuration is needed to identify a BoQ speaker's role: a QUIC client, a QUIC server, or any (Don't care). The default value is "any".
- A new "BGP over QUIC" capability is defined below to signal whether the BoQ speaker is a QUIC client, a QUIC server, or any (Don't care).

Detail and conflict resolution are defined in the draft.

To-dos on the Roadmap

- Finite State Machine
- Graceful Restart
- Connection Migrations
- Channel Prioritization
- Operational Considerations for TLS certificate
- ...

Next Steps

- Reviews and comments are welcome
- GitHub Repository: github.com

Thanks!