QUIC Updates

ATS Fall 2017 Summit

Masaori Koshiba <masaori@apache.org>
Masakazu Kitajo <maskit@apache.org>

Agenda

- IETF-QUIC
- Project Status
- QUIC Interop Session Report
- Development Notes
- Open Questions
- Future Plans

Google-QUIC ≠ IETF-QUIC

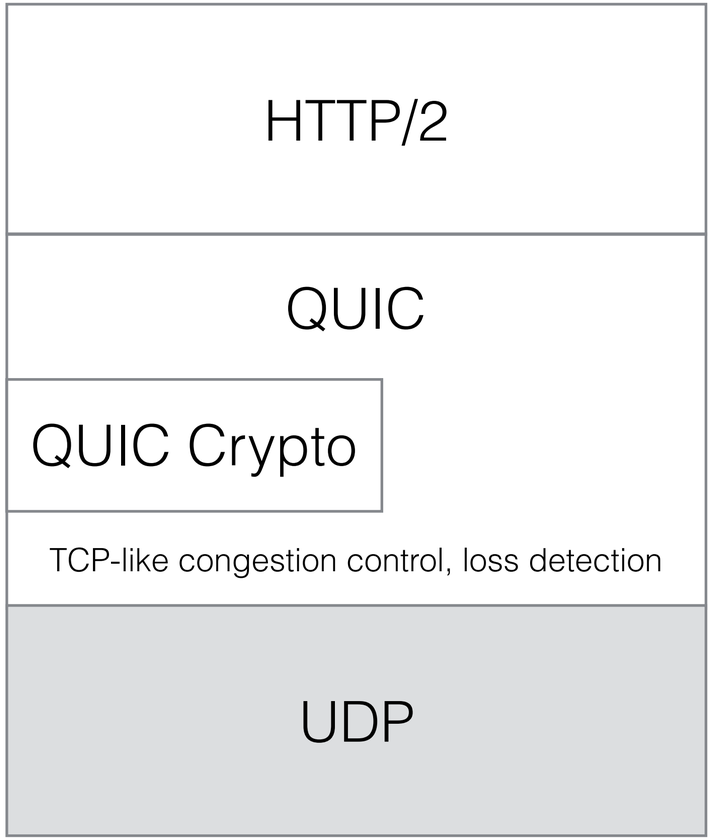

Google-QUIC

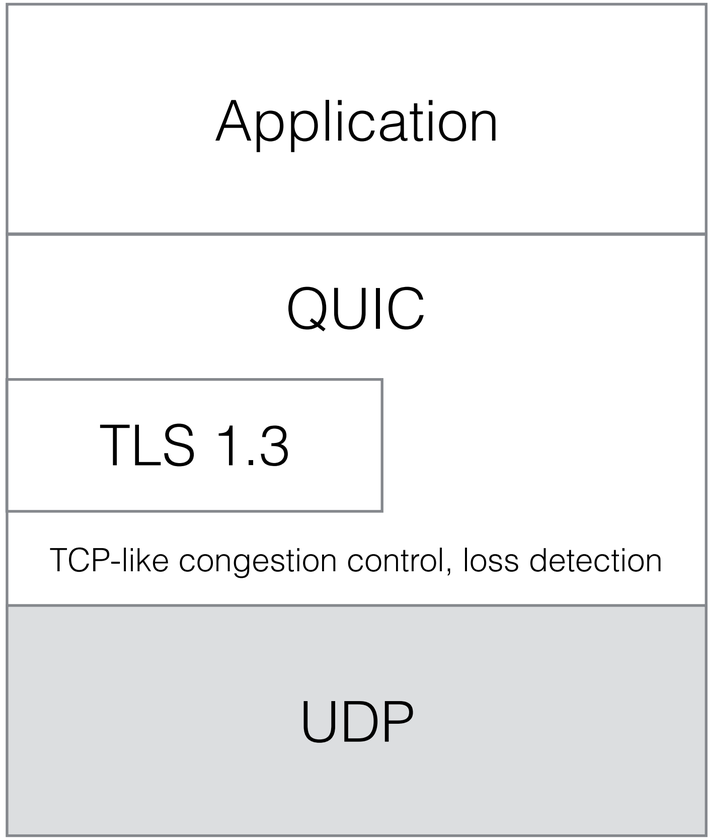

IETF-QUIC

IETF-QUIC
≒ Improved TCP
+ TLS 1.3
+ Improved HTTP/2
(as Application)

The QUIC Spec is changing FAST

- The 4th draft series was released on Jun. 13
- The 5th draft series was released on Aug. 15 (implementation draft)
- The 6th draft series was released on Sep. 22
- The 7th draft series was released on Oct. 13

Lots of things are still under discussion…

Last few months of QUIC on ATS

- Announced QUIC development project - *5 months ago*
- Open the initial code #2342 - *2 months ago*
  - Made it work on Fedora and FreeBSD
  - Fixed many things
  - Implemented some QUIC features
- Attended Interop Session - *3 weeks ago*

QUIC Implementation Targets

First Target

- Handshake (1-RTT)
- Version Negotiation
- Data transfer (encrypted)
- Connection Close

Seconds Target

- HTTP/0.9
- Server Stateless Retry
- Stateless Reset
- Flow Control

Third Target (?)

- Handshake (0-RTT)
- Congestion Control
- PMTU Discovery
- Key Updates

...

QUIC Implementation Targets

First Target

- Handshake (1-RTT)
- Version Negotiation
- Data transfer (encrypted)
- Connection Close

Seconds Target

- HTTP/0.9
- Server Stateless Retry
- Stateless Reset
- Flow Control

Third Target (?)

- dshake (0-RTT)
- Congestion Control
- PMTU Discovery
- Key Updates

…

Implemented targets on ATS

First Target

- ~~Handshake (1-RTT)~~
- ~~Version Negotiation~~
- ~~Data transfer (encrypted)~~
- ~~Connection Close~~

Seconds Target

- ~~HTTP/0.9~~
- Server Stateless Retry
- ~~Stateless Reset~~
- ~~Flow Control~~

Third Target (?)

- Handshake (0-RTT)
- Congestion Control
- PMTU Discovery
- Key Updates

...

Project Status

- Progress: ✔ ***Going well***
  - Most items for the last interop are implemented
  - Catching up the latest changes
- Contributions:
  - **175** commits (+16,264 –55 lines)
  - 6 committers (**+4**), 2 contributors (**+2**)

QUIC WG Interop Session

- Implementations check connectivity to each other
- Server Side: publish endpoint
  - e.g. public endpoint of ATS (quic.ogre.com:4433)
- Client Side: access to the endpoint

QUIC WG Interop Session

Oct. 3 at Seattle

client ↓ | server →	Minq	Mozquic	Quicly	Quant	ngtcp2	NGINX+ngtcp2	mvfst	picoQUIC	WinQUIC	ATS (HNDC9!)	AppleQUIC	f5
Minq	HDC9F	HD9RF	HD9		HD9		H		HD9	H		
Mozquic	HNDC9F	HNDC9R!F	HDC9RF	H	HNDC9F		HND		HND9F	HNDC		
Quicly	HD9	HD9FR	HD9RF		HD9		H	HDC9	HD9	HD9		
Quant				HDC			H					
ngtcp2	HNDC9	HNDC9R	HD9R		HNDC9F		HN		HNDC9	HNDC9		
mvfst		HN	HND	H			HND9F					
picoQUIC	HDC9	HDC9F	HD9F	HDC	HDC9			HNDC9R!F	HDC9F			
WinQUIC	HNDC9F	HNDC9!F	HDC9	HC	HNDC9F		HNDC9		HNDC9F0	HNDC9		
AppleQUIC	H	H		H	H		H		H	HND	HNDCF9	
f5												

	Legend:				
	self-test	interop	known broken	unknown	N/A

To Test:	H = Handshake	
	N = Version Negotiation	
	D = Stream Data (encrypted)	
	C = Close	
	9 = HTTP/0.9 exchange	-- can we agree on what this? Perhaps 3 parallel streams?
	R = server stateless retry	
	! = stateless reset	
	F = flow control	does this cover advertising and enforcing flow control?
	0 = 0-RTT	

QUIC WG Interop Session

Oct. 3 at Seattle

Feature	Result	Notes
Handshake	6 / 6	
Version Negotiation	4 / 6	A client didn’t support ALPN
Data transfer (encrypted)	5 / 6	
Connection Close	3 / 6	
HTTP/0.9	3 / 6	A client didn’t support HTTP/0.9 Found problem with Flow Control
Server Stateless Retry	N/A	
Stateless Reset	0 / 6	No implementation tested
Flow Control	0 / 6	Found bug (#2620)

QUIC WG Interop Session

Oct. 3 at Seattle

Found many issues

- Crashes
  - #2604, #2606, #2607, #2608, #2610, #2611, #2613, #2619
- Bugs / Improvements
  - #2609, #2616, #2620

(Search "is:issue label:QUIC created:2017-10-03..2017-10-05")

Development Notes : TLS 1.3

- OpenSSL 1.1.1-dev (TLS 1.3 draft-21)

```
$ ./config enable-tls1_3
```

- records.config

```
CONFIG proxy.config.ssl.server.cipher_suite STRING
  TLS13-AES-128-GCM-SHA256:
  TLS13-AES-128-GCM-SHA256:
  TLS13-CHACHA20-POLY1305-SHA256:…
```

https://www.openssl.org/blog/blog/2017/05/04/tlsv1.3/

Development Notes : Configs

- records.config

```
CONFIG proxy.config.udp.threads INT 1
CONFIG proxy.config.http.server_ports STRING 4433:quic

# optional
CONFIG proxy.config.quic.no_activity_timeout_in INT 30
```

- ssl_multicert.config

```
dest_ip=* ssl_cert_name=/CERT ssl_key_name=/PRIVATE/KEY/OF/CERT
```

Development Notes : C++11

- smart pointer, enum class, constexpr and override
- Unfriendliness of lib/ts containers to std::unique_ptr
  - e.g. Map / DLL / Queue
- std::swap v.s. swap in lib/ts/Vec.h (#2357)
- std::make_unique (c++14) (c4c3bd9c)
- type alias

```
using QUICStreamFrameUPtr =
std::unique_ptr<QUICStreamFrame, QUICFrameDeleterFunc>;
```

Development Notes

- #pragma once
- Placement new with ClassAllocator (alternative to init())

```
MyClass *buf = myClassAllocator.alloc();
buf->init(arg1, arg2)
```

v.s.

```
MyClass *buf = myClassAllocator.alloc();
new (buf)MyClass(arg1, arg2);
```

Development Notes : Unit Tests

Catch & EventProcessor

```
struct EventProcessorListener : Catch::TestEventListenerBase {
  using TestEventListenerBase::TestEventListenerBase;

  virtual void
  testRunStarting(Catch::TestRunInfo const &testRunInfo) override
  {
    …
    eventProcessor.start(TEST_THREADS);
    …
  }
};
CATCH_REGISTER_LISTENER(EventProcessorListener);
```

Development Notes : Abstract Classes

- Abstract Calss QUICNetVConnection has many features
- Divide them in small pieces

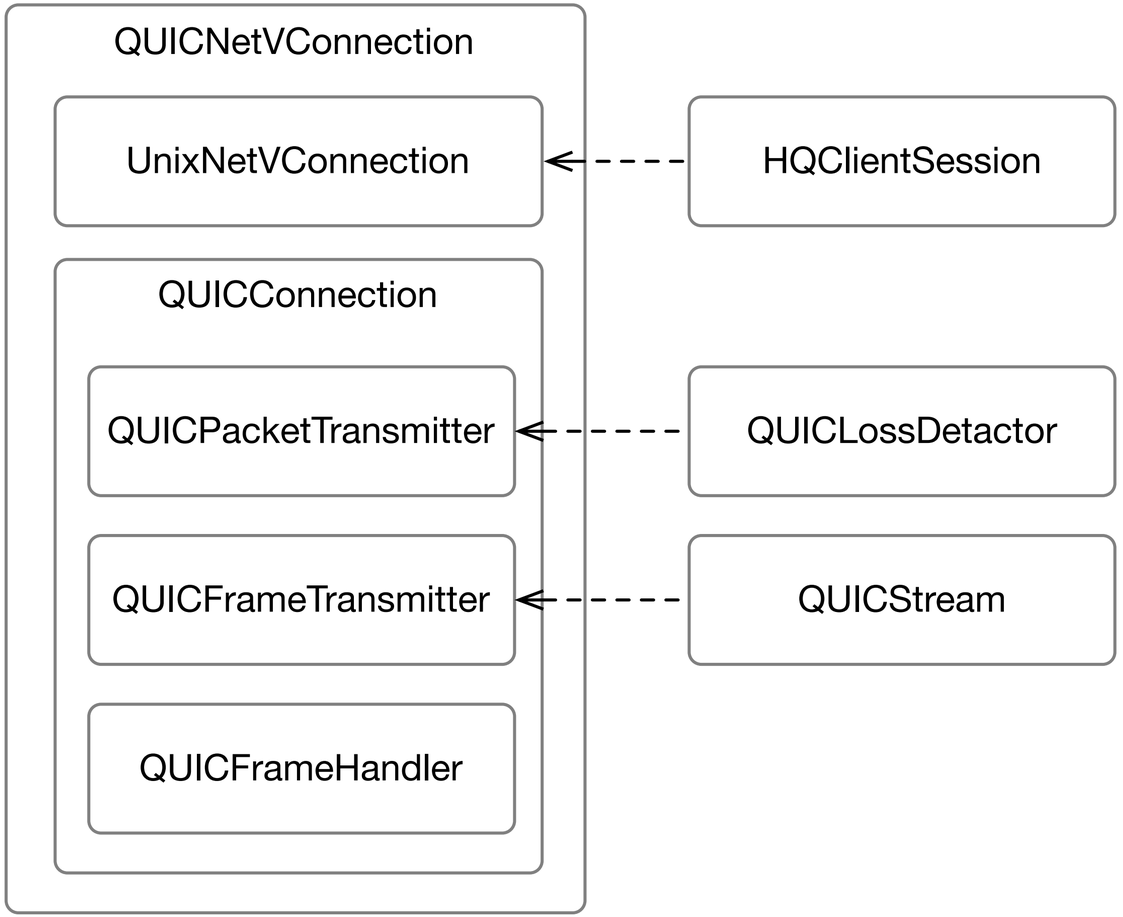

Development Notes : Branches

QUIC Development Branches

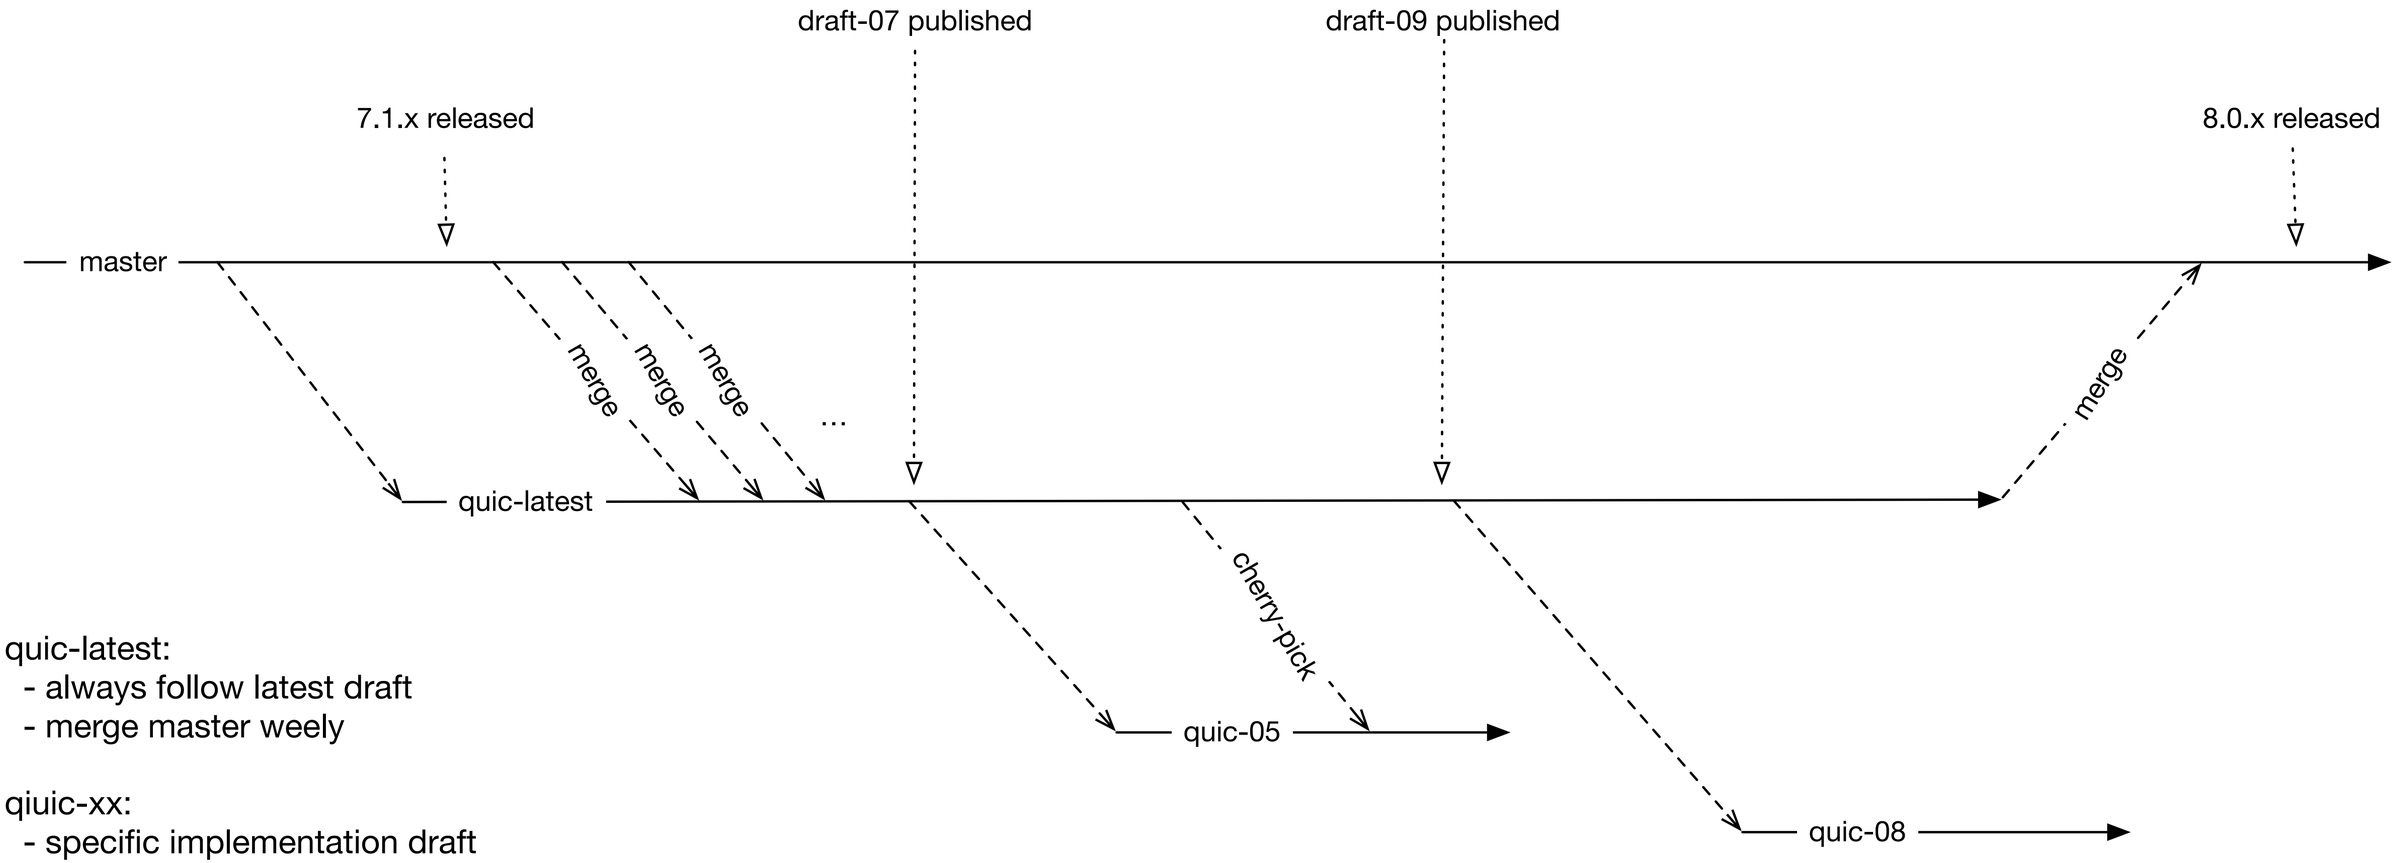

quic-latest:
- always follow latest draft
- merge master weely

qiuic-xx:
- specific implementation draft
- for interop

ATS Integrations

Areas we haven’t touched

- Stats (with consistent naming convention)
- Logs
- Plugin APIs and Hooks
- AUTest
- Origin server side QUIC connection (client implementation)

Known Issues

- Stateless Server Retry
  - TLS 1.3 HelloRetryRequest cookie
  - OpenSSL doesn’t support yet (openssl/#4283, openssl/#4435)

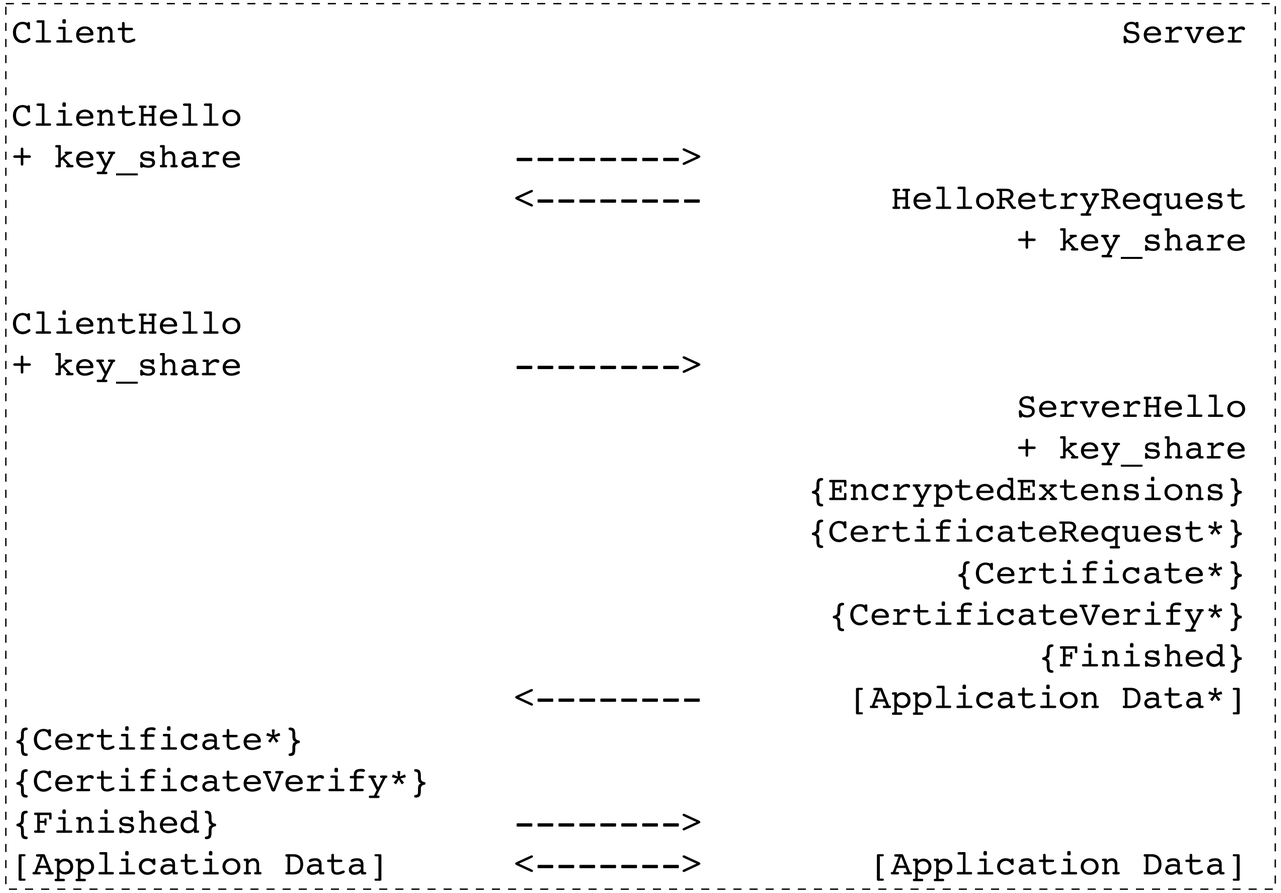

*https://tools.ietf.org/html/draft-ietf-tls-tls13-21#section-2.1

Open Questions

- Performance
  - Not measured since there’s no tools
- Concurrency Containers (Map, Queue) (#2478, #2555)
- Design of QUICNetVConnection (#2558)
- Design of QUICStream / HQClientTransaction

Concurrency Containers #2478, #2555

Discussion Session in Day 2

- Use Case 1 : Finding QUICNetVConnection

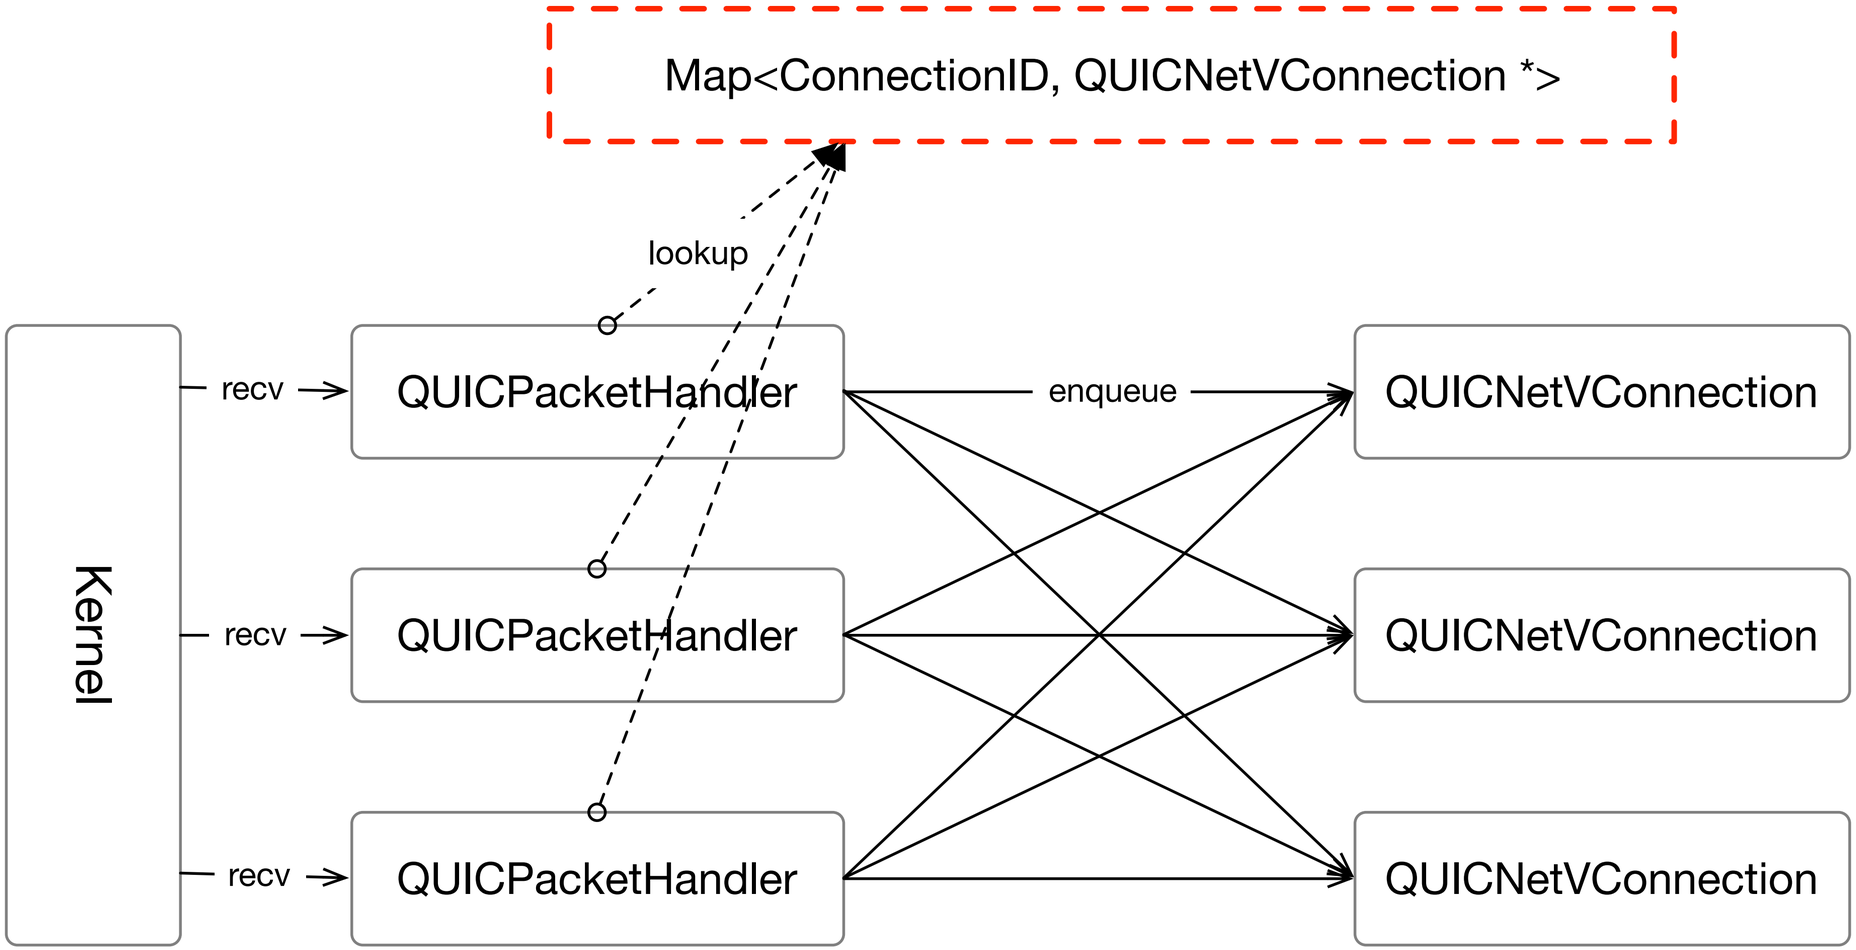

Concurrency Containers #2478, #2555

Discussion Session in Day 2

- Use Case 2 :Enqueue QUICPacket

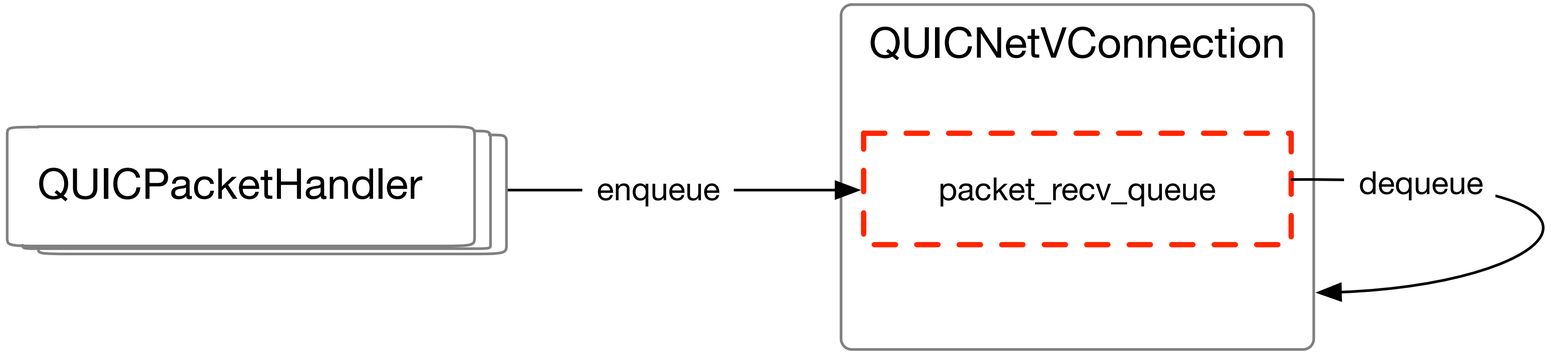

- Use Case 3 :Enqueue QUICFrame

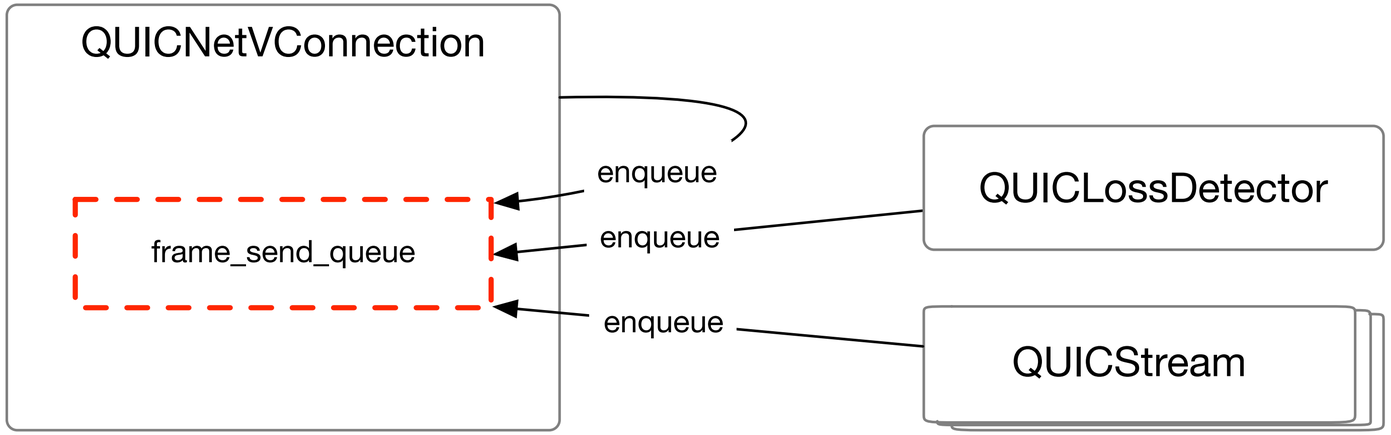

Design of QUICNetVConnection #2558

- Problems
  - QUICNetVConnection inherits UnixNetVConnection for InactivityCop
  - UnixNetVConnection is designed for TCP
  - QUICNetVConnection is not using read / write operation of UnixNetVConnection (net_read_io / load_buffer_and_write)
- Ideas
  - QUICNetVConnection inherits NetVConnection directly
  - Add QUICNetHandler (QUIC version of NetHandler)

Design of Stream

- Common Point
  - Own read_vio / write_vio
- Unique Point
  - Chunked Encoding

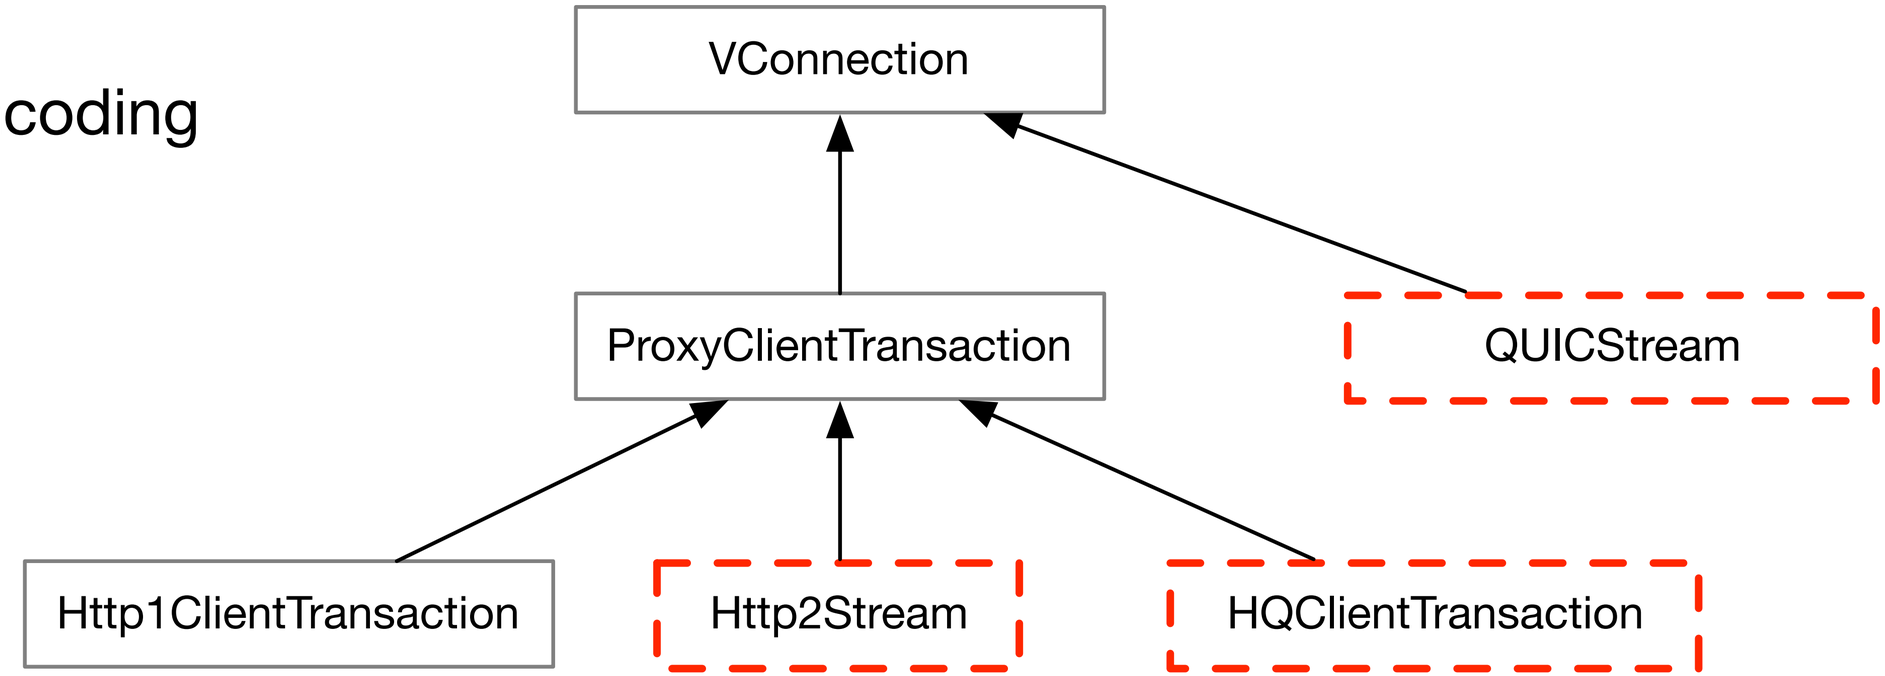

Design Overview

HTTP/1.1 over TCP

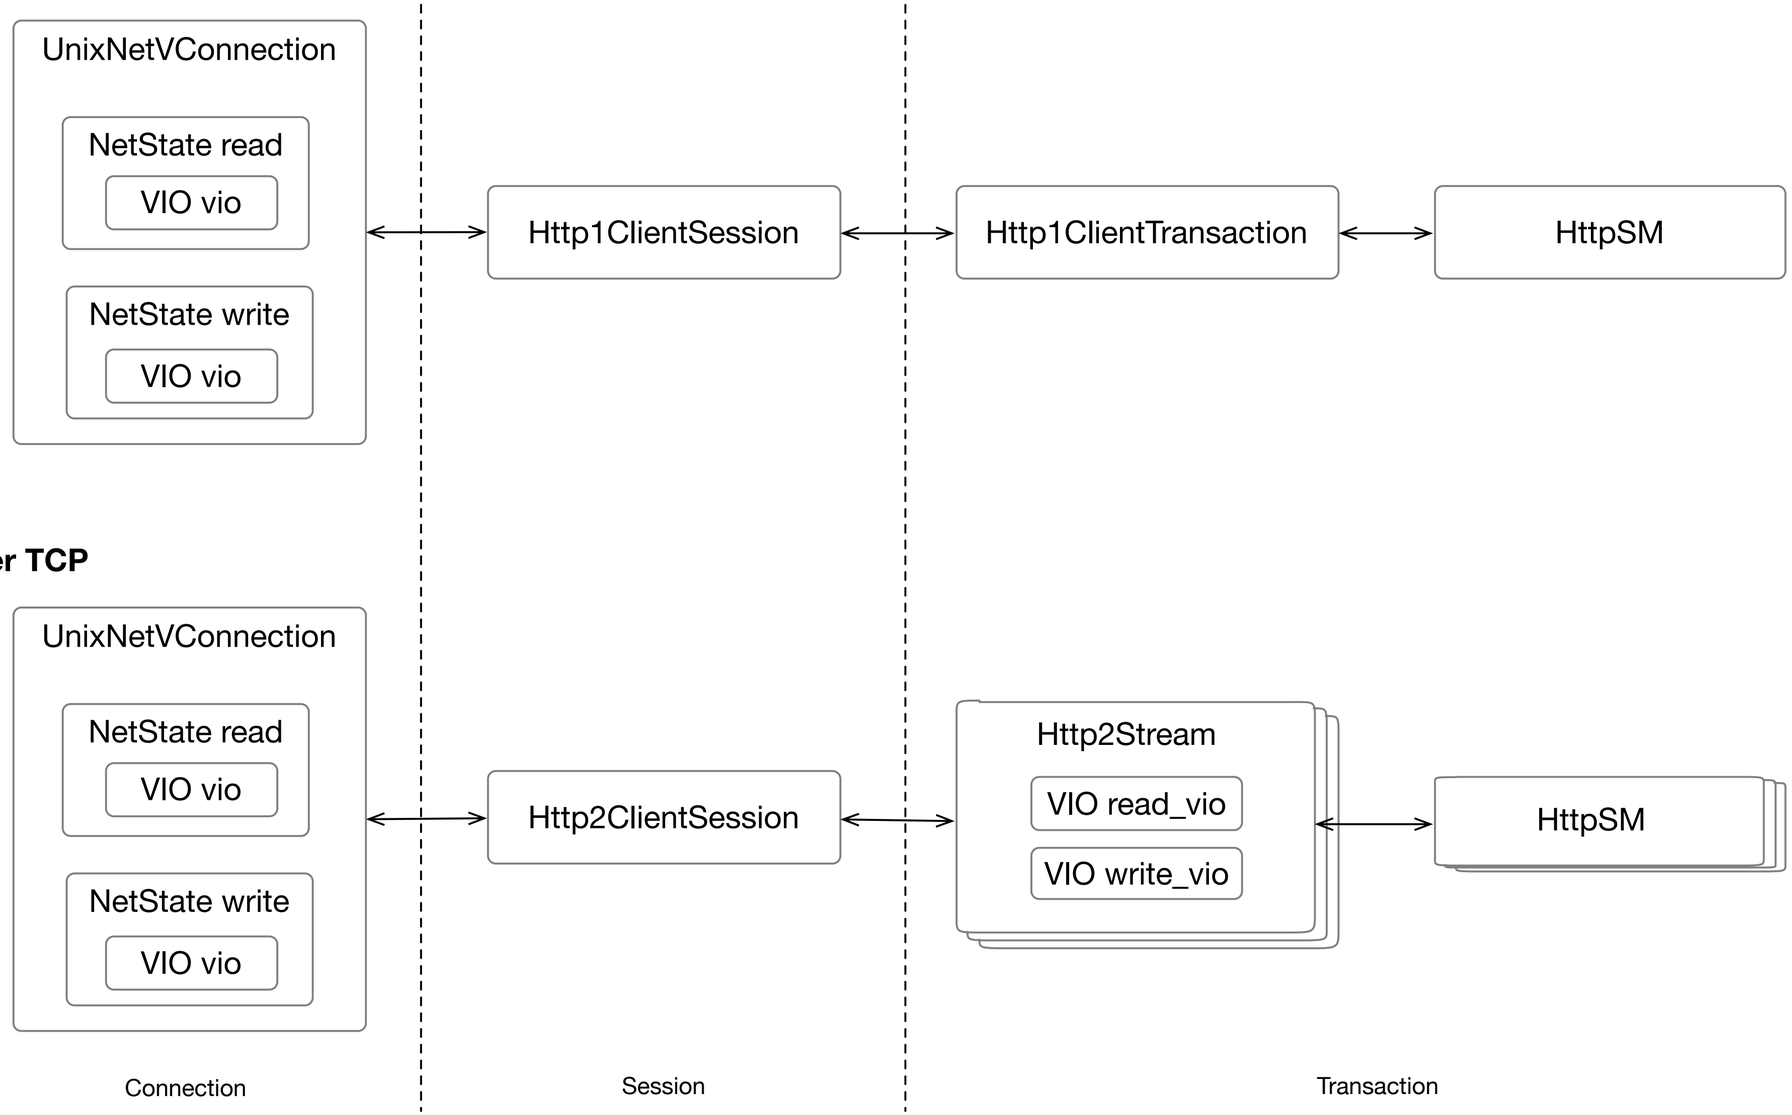

HTTP/2 over TCP

Design Overview

HTTP over QUIC

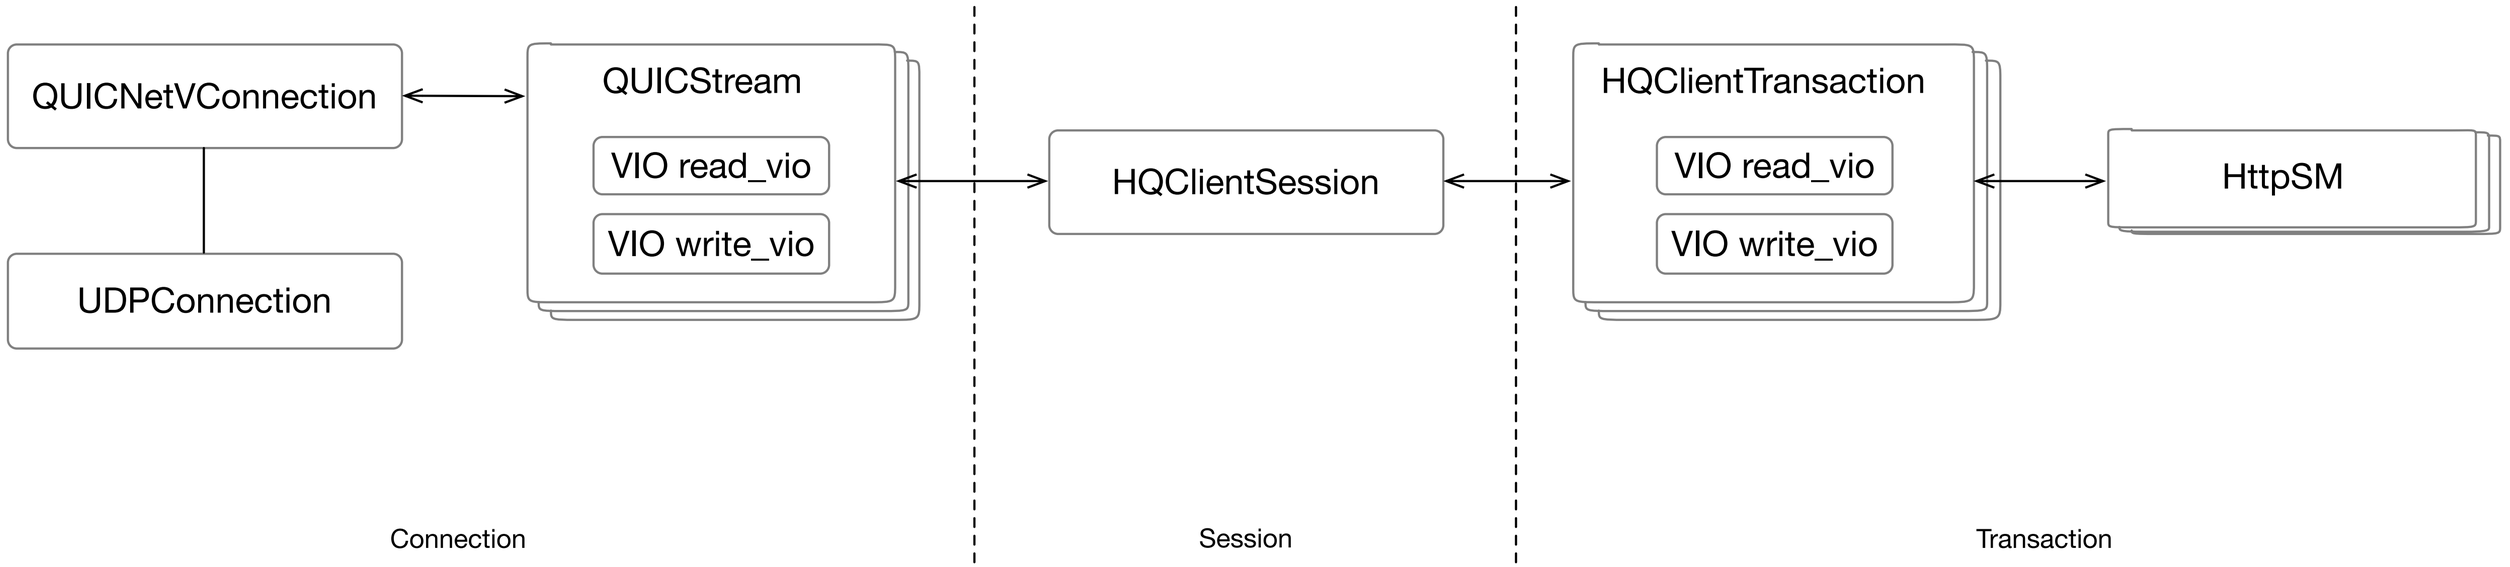

Future Plans

Date	Event	Note
2017 Nov.		• Implement more QUIC features • Start integrating stats, logs, etc.
2018 Jan.	• ATS QUIC Hackathon in Japan • QUIC Interop in Australia	• Merge quic-latest to the master (?)
2018 Apr. (?)	• ATS Spring Summit	• Report the result of interop
2018 Jul. (?)	• ATS 8.0 release	• Experimental QUIC support (?)

Thanks

- Documentations & Links
  - https://cwiki.apache.org/confluence/display/TS/QUIC
- Code
  - https://github.com/apache/trafficserver/tree/quic-latest
- Questions?
  - List: dev@trafficserver.apache.org
  - IRC: #traffic-server on freenode

HTTP/0.9 over QUIC