

QUIC Updates

ATS Euro-Tour 2018

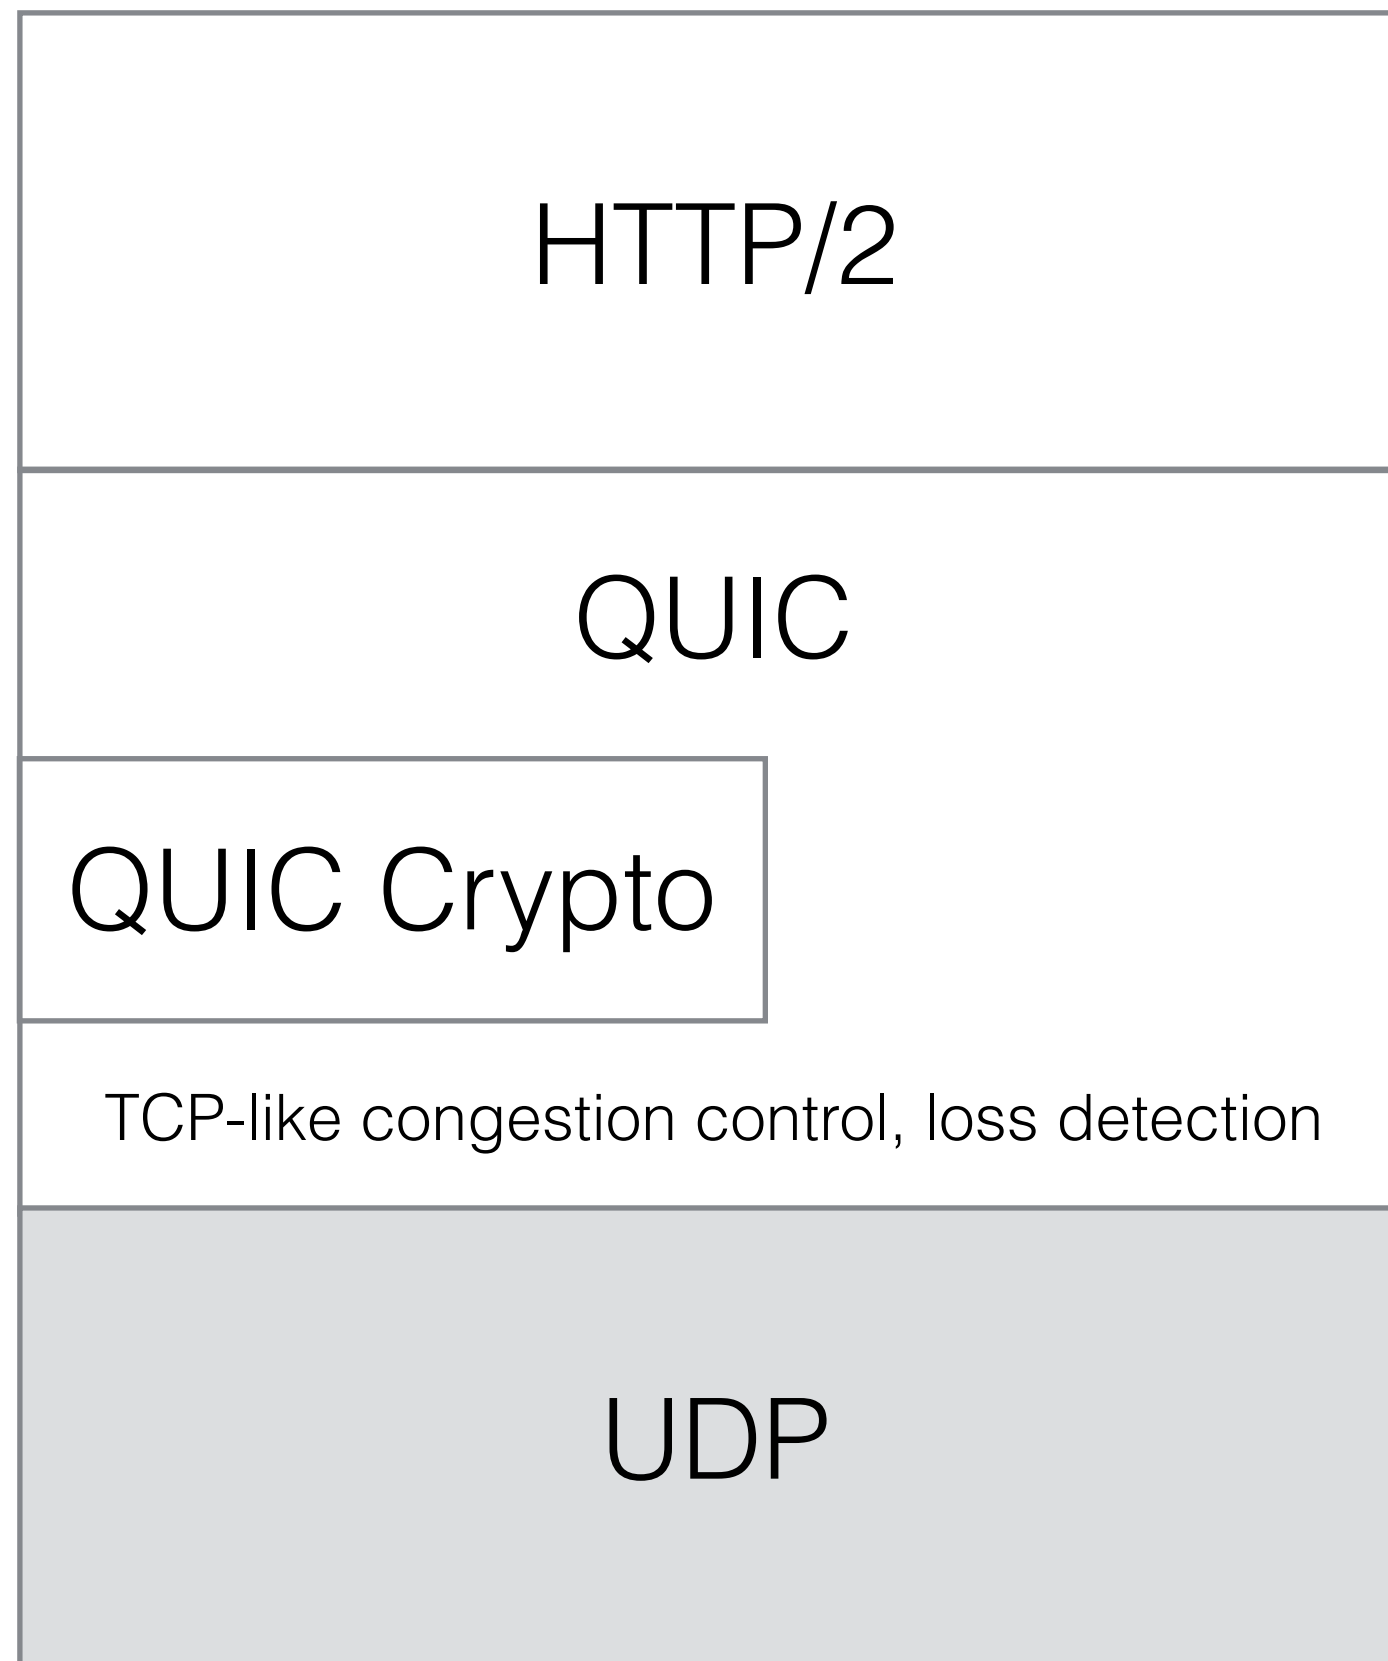
Masakazu Kitajo <maskit@apache.org>
Masaori Koshiba <masaori@apache.org>



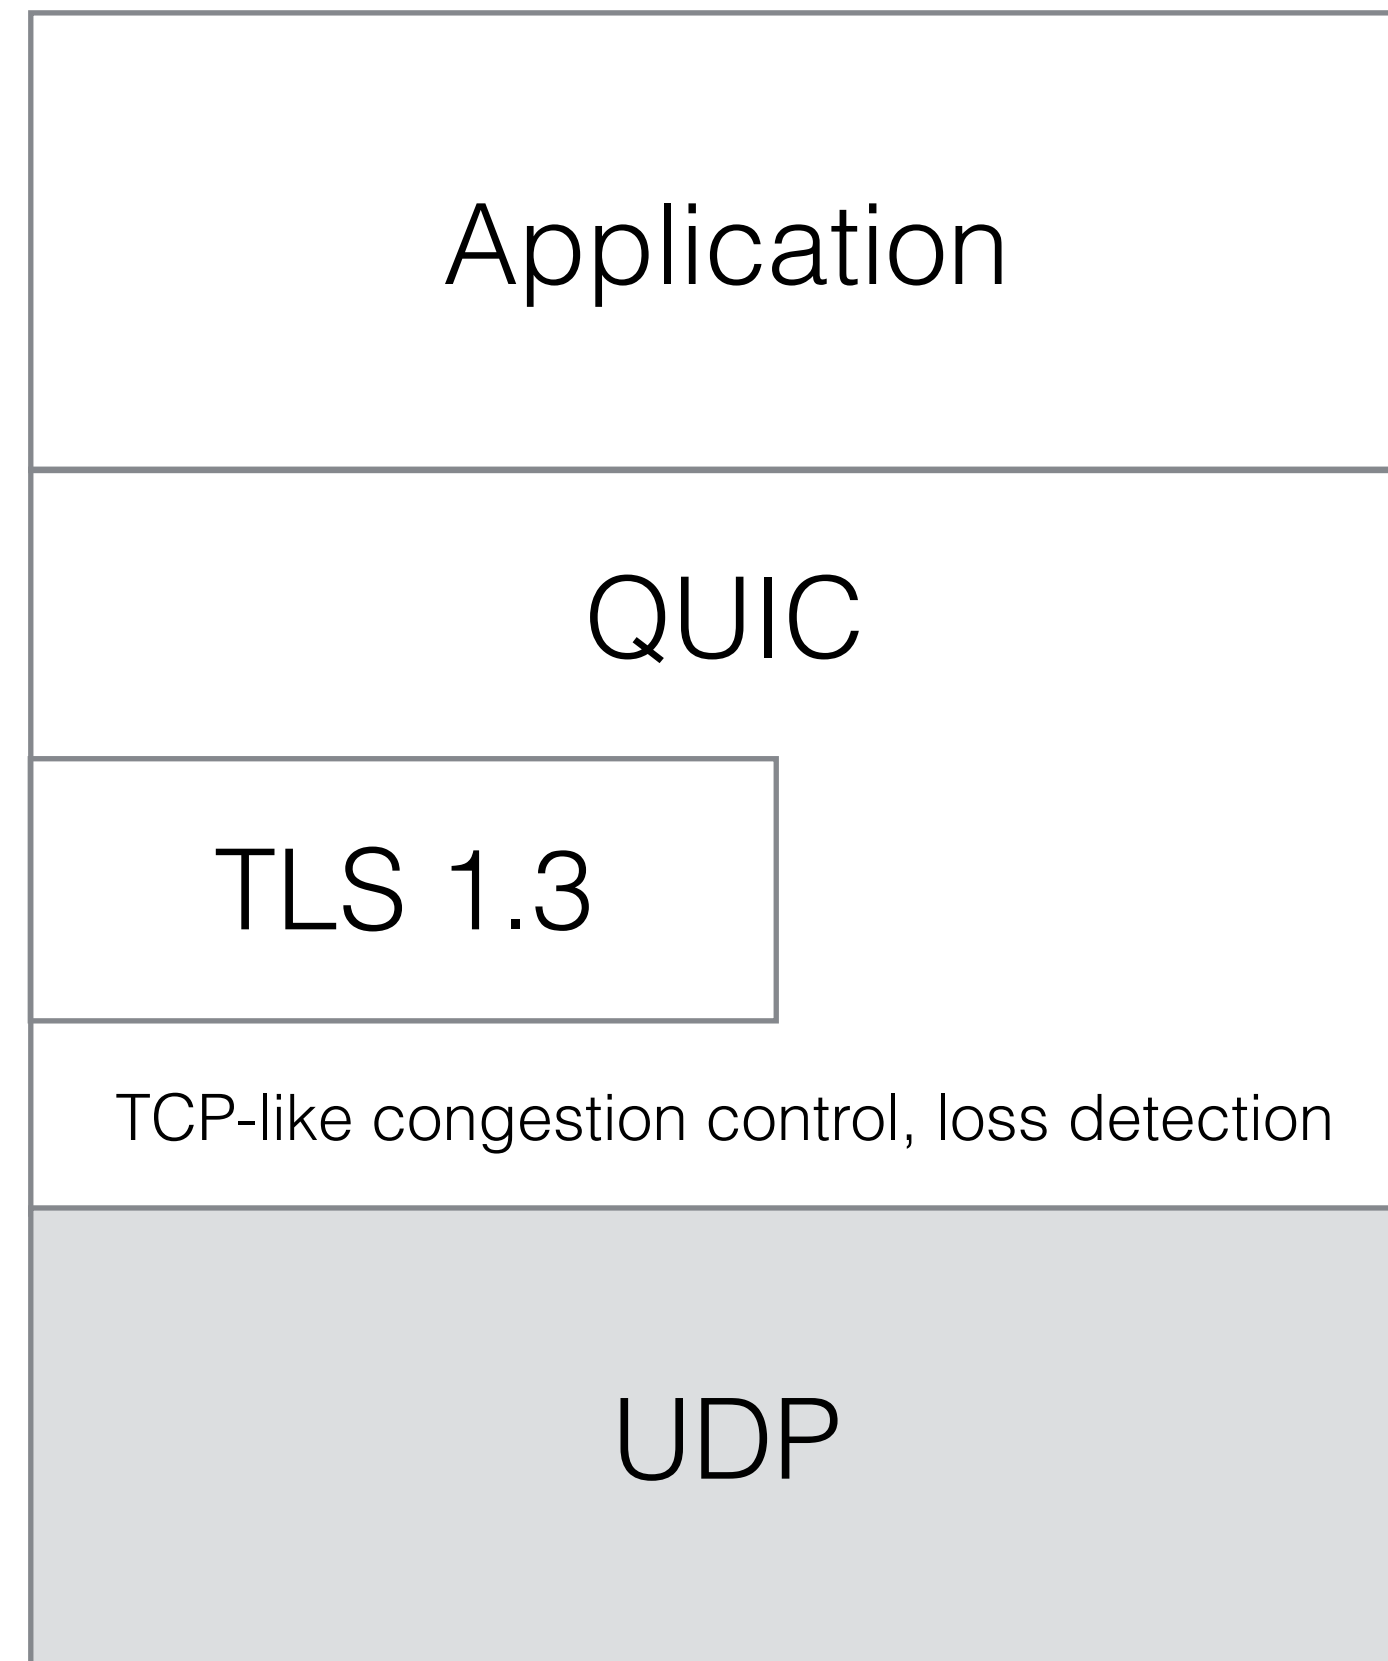
Agenda

- IETF-QUIC
- Project Status
- QUIC Interop Session Report
- Development Notes
- Known Issues
- Future Plans

Google-QUIC \neq IETF-QUIC



Google-QUIC



IETF-QUIC

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



IETF-QUIC

- The 7th draft series was released on Oct. 13
- The 8th draft series was released on Dec. 5
- The 9th draft series was released on Jan. 28
- The 10th draft series was released on Mar. 5
- The 11th draft series was released on Apr. 17

Nov 2018	Header Compression for HTTP over QUIC to IESG draft-ietf-quir-qcram
Nov 2018	QUIC Applicability and Manageability Statement to IESG draft-ietf-quir-applicability draft-ietf-quir-manageability
Nov 2018	HTTP/2 mapping document to IESG draft-ietf-quir-http
Nov 2018	TLS 1.3 Mapping document to IESG draft-ietf-quir-tls
Nov 2018	Loss detection and Congestion Control document to IESG draft-ietf-quir-recovery
Nov 2018	Core Protocol document to IESG draft-ietf-quir-transport
Sep 2018	Version-Independent Properties of QUIC to IESG draft-thomson-quir-invariants

November? Seriously?

~~Lots of~~ Several things are still under discussion...



TLS 1.3

- IESG approved the document (draft-28)
- OpenSSL 1.1.1 supports it
- Cipher suite need to be updated

```
CONFIG proxy.config.ssl.server.cipher_suite STRING
  TLS13-AES128-GCM-SHA256:
  TLS13-AES256-GCM-SHA384:
  TLS13-CHACHA20-POLY1305-SHA256:...
```

ATS QUIC Project: Recent Activity

- Joined Interop (online) - 5 months ago
 - Checked interoperability with draft-08 impl.
- ATS QUIC Hackathon - 4 months ago
- Joined Interop - 4 months ago
 - Checked interoperability with draft-08 impl.
- Joined Interop (online) - 2 months ago
 - Checked interoperability with draft-09 impl.

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

- Resumption
- Handshake (0-RTT)
- Loss Recovery

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

- ~~Resumption~~
- ~~Handshake (0-RTT)~~
- ~~Loss Recovery~~

ATS QUIC Project: Current status

- Progress: ✓ *Going well*
 - ALL items for the last interop are implemented (Apr. 17th)
 - Catching up the latest changes (draft-11)
- Contributions:
 - **250+** commits on quic-latest branch (**+27,589** **-144** lines)
 - 7 committers (**+1**), 1 contributors (**+1**)

QUIC WG Interop Session

- Server side: publish endpoints
- Client side: access to the endpoints and sends QUIC packets



QUIC WG Interop Session

Mar. 20 at IETF101

client ↓ server →	Minq	Mozquic	Quicly	quant	ngtcp2	mvfst	picoQUIC	winquic	f5	ATS	Pandora	AppleQUIC	ngx_quic	!gquic	quic-go	quicker
Minq	-	HS	HS	-	-	HS	-	-	-	-	-	-	-	-	-	-
Mozquic	-		VHDCRZS	VHDCRZS	VHDCRZ	VHDCRZ	VHDCRZS	VHDCRZS	VHC	VHDCR	VHC	-	VH	-	-	-
Quicly	-	VHDRZS		VHDRZS	VHDRZ	VHDCRZ	VHDRZS	VHDRZS	VHD	HDRZ		-		-	-	-
quant	-	VHDCRZS	VHDCRZS	VHDCRZS	VHDCRZ	VHDCRZ	VHDCRZS	VHDCRZS	VHDC	VHDCRZ	VHDC	-	VHC	-	-	-
ngtcp2	-	VHDCRZS	VHDCRZS	VHDCRZS	VHDCRZ	VHDCRZ	VHDCRZS	VHDCRZS	VHC	VHDCR	VHDC	-	H	-	-	-
mvfst	-	VHDCR		VHDCR	VHDCR	VHDCR	VHDCR	VHDCR	VHC	VHDCR	VHDC	-		-	-	-
picoQUIC	-	VHDCRZS	VHDCRZS	VHDCRZS	VHDCRZ	VHDC	VHDCRZS	VHDCRZS	VHC	VHDCR	VHDC	-	H	-	-	-
winquic	-	VHDCRZS	VHDCRZS	VHDCRZS	VHDCRZ	VHDCRZ	VHDCRZS	VHDCRZS	VHDC	VHDCR	VHDC	-	VHC	-	-	VHDCRZ
f5	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-	-
ATS	-	H	H	H	H	H	X	H	H	H	H	-	-	-	-	-
Pandora	-	VHDC		VHDC	VHDC	VHDC	VHDC	VHDC	VHDC	VHDC	VHDC	-		-	-	-
AppleQUIC	-	VHX		VHDC	VHD	VHDC	VHDS	VHD	VX	VHD	VHX			-	-	-
ngx_quic	-													-	-	-
!gquic	-													-	-	-
quic-go	-													-	-	-
quicker	-	VHDCRZ			VHDCRZ		VHDC	VHDCRZ		VHDCRZ				-	-	-

Version Negotiation	V
Handshake	H
Stream Data	D
Connection Close	C
Resumption	R
0-RTT	Z
Stateless Retry	S
N/A	-
Known broken	X
Self-test	



QUIC WG Interop Session

Mar. 20 at IETF101

Feature	Result	Notes
Handshake	10 / 10	
Version Negotiation	9 / 10	Small issues found
Data transfer (encrypted)	10 / 10	
Connection Close	8 / 10	Small issues found
HTTP/0.9	OK	No issues found
Server Stateless Retry	N/A	Not ready to test
Stateless Reset	N/A	No test scenarios
Flow Control	OK	No issues found
Resumption	8 / 10	No issues found (2 clients haven't supported it yet)
0-RTT	3 / 10	Issues found
Loss Recovery	OK	No issues found

Development Notes

- Concurrent Containers
- Server side connection
- WRITE_READY event loop
- QUICNet/QUICNetProcessor

Development Notes : Concurrent Containers

- Table
 - Purpose: Connection ID -> QUICNetVConnection
 - Implementation: MT_Hashtable in ts/MT_Hashtable.h
- Queue
 - Purpose: Packet send/receive queue
 - Implementation: Queue in ts/List.h (to be replaced with ASLL)

Related session: Hash tables and multi-threading on Thursday

Development Notes : Server side connection

- QUIC
 - QUICNetVConnection
- HTTP/QUIC
 - HQServerSession ?
 - HQServerTransaction ?
- Protocol Negotiation
 - HostDB ?

Related session: HTTP/2 to origin on Wednesday

Development Notes : WRITE_READY event loop

- QUICNetVConnection schedules a WRITE_READY event for itself

```
case QUIC_EVENT_PACKET_WRITE_READY: {  
    this->_close_packet_write_ready(data);  
    error = this->_state_common_send_packet();  
    // Reschedule WRITE_READY  
    this->_schedule_packet_write_ready(true);  
    break;  
}
```

```
uint32_t  
QUICNetVConnection::transmit_packet(QUICPacketUPtr packet)  
{  
    uint32_t npackets = this->_transmit_packet(std::move(packet));  
    this->_schedule_packet_write_ready();  
    return npackets;  
}
```


Development Notes : QUICPollEvent/Cont

- QUICPollEvent
- QUICPollCont
- QUICNet
- QUICNetProcessor

Known Issues

- Low performance (PR#3552)
 - 3.6kbsp -> 1Mbps -> 8Mbps
 - Getting better...
- CI for QUIC stuff (PR#3551)
 - QUIC requires TLS 1.3 (OpenSSL 1.1.1)

Future Plans

Date	Event	Note
2018 Jun.	<ul style="list-style-type: none">• ATS 8.0 release• QUIC Interop	<ul style="list-style-type: none">• Any QUIC stuff won't be included in the next release
2018 Sep. (?)	<ul style="list-style-type: none">• QUIC Interop	
2018 Oct. (?)	<ul style="list-style-type: none">• ATS Fall Summit	<ul style="list-style-type: none">• Report the interop result
2018 Nov.	<ul style="list-style-type: none">• QUIC WG Last Call	<ul style="list-style-type: none">• Merge quic-latest to the master (?)
2019 (?)	<ul style="list-style-type: none">• ATS 8.1 release (?)	<ul style="list-style-type: none">• 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

