

Pluginizing QUIC

Quentin De Coninck, **François Michel**, Maxime Piraux, Florentin Rochet, Thomas Given-Wilson, Axel Legay, Olivier Pereira, Olivier Bonaventure

SIGCOMM 2019 UCLouvain, Belgium



QUIC

- Currently under standardization at IETF
- Providing TCP/TLS1.3 service atop UDP
 - (Nearly) everything is encrypted
 - Reliable, in-order, multistream data transfer
- Frame-based protocol
- Large companies are deploying both client and server sides

F

Н

• Limited interoperability today



F

F



Application Requirements Evolve



Simple reliable file transfer

Mobility, (live) video streaming

Latency critical, partially unreliable

Transport protocols need to adapt!





What If We Could Have...

A base, simple implementation enabling

Personalized Transport Functionalities Per-Connection Deployment



Agenda

- Motivations
- Pluginized QUIC Design
- Evaluating PQUIC with Use Cases
- Conclusion and Future Works

Pluginized QUIC



- Revisit the structure of protocol implementations
- Transport protocol = set of basic functions (protocol operations)







Revisiting Protocol Implementation Design

What do we need?

- Running plugins in a safe environment
- Identifying protocol operations
- Providing an API to the protocol plugins

Running Plugins

- Plugin = set of bytecodes
 - 1 bytecode \rightarrow 1 function
 - Hardware/OS independent
 - Kept under control
- Rely on a virtual machine
 - Isolate bytecode from host implementation





eBPF Virtual Machine

• Lightweight, integrated in Linux kernel since 2014

- RISC instruction set (~100)
 - ALU, memory and branch purposes
- Bytecode recompiled to native architecture
- Verifier to ensure programs safety
- Dedicated, isolated stack memory
 - But no persistence
- Rely on a user-space implementation
 - With relaxed verifier
 - With persistent heap memory



Identifying Protocol Operations



Modifying/Adding Protocol Operations



Changing the Retransmission Timer



Inserting Monitoring Plugin



Linking Plugins to the Core Implementation

- Exposing connection fields
- Offering persistent heap memory
- Retrieving data in plugin memory
- Calling other protocol operations





Communication between Plugin Bytecodes



Agenda

- Motivations
- Pluginized QUIC Design
- Evaluating PQUIC with Use Cases
- Conclusion and Future Works

Evaluation Overview

Implementation based on picoquic

- IETF QUIC in C language
- Evaluation in our lab
 - Use NetEm and HTB to configure links
- Experimental design
 - 139 points
 - Consider median over 9 runs



Metric	Minimum	Maximum
One-way delay (ms)	2.5	25
Bandwidth (Mbps)	5	50

Exploring Use Cases

- Monitoring
- A QUIC VPN
- Multipath
- Forward Erasure Correction

Plugin	Lines of C Code	Number of bytecodes
Monitoring	500	14
QUIC VPN	500	11
Multipath	2600	32
Forward Erasure Correction	2500	51

Exploring Use Cases

- Monitoring
- A QUIC VPN
- Multipath

Read paper for other use cases

Forward Erasure Correction

Plugin	Lines of C Code	Number of bytecodes
Monitoring	500	14
QUIC VPN	500	11
Multipath	2600	32
Forward Erasure Correction	2500	51

1st Plugin: QUIC VPN

Alternative to DTLS, IPsec



- → But QUIC provides **reliable**, **in-order** data delivery
 - Datagram plugin
 - Unreliable data delivery
 - Application API
 - Protocol operation



QUIC VPN in Action

Compare download time for a single file transfer using TCP_{Cubic} over 1 path





2nd Plugin: Multipath QUIC



Implement the IETF Multipath draft

- Link paths to specific networks
 - Connection not bound to 4-tuple
- Exchange IP addresses
- Create paths over pair of IP addresses
 - Full-mesh
- Schedule packets over paths
 - Round-robin fashion



Multipath QUIC in Action

Speedup Ratio =

Download Completion Time Single Path QUIC Download Completion Time Multipath QUIC

GET request over a dedicated stream

• Consider 2-path network with **symmetric characteristics**



Combining Plugins into a Multipath QUIC VPN







Combining Plugins into a Multipath QUIC VPN



Combining Plugins into a Multipath QUIC VPN



Multipath QUIC VPN in Action

Compare download time for a single file transfer using TCP_{Cubic}



Agenda

- Motivations
- Pluginized QUIC Design
- Evaluating PQUIC with Use Cases
- Conclusion and Future Works

Conclusion and Future Works

- PQUIC: Dynamically extends protocol implementations
- Common **pluginizable client**, **specialized server** implementation
 - Possible through secure plugin exchange
- Four different use-cases fully implemented with plugins
 - Monitoring, VPN, Multipath, Forward Erasure Correction
- Future works
 - Approach applicable to many other protocols
 - How can we have independent, interoperable PQUIC implementations?
 - Specification of the VM, protocol operations, API,...

https://pquic.org

01011 10010

Thanks for your attention!

https://pquic.org

Extending/Customizing TCP Is Hard

- TCP Option Negotiation
- Extension deployment challenges
 - Specification process (IETF)
 - Implementation puzzle
 - Client waits for server support
 - Server waits for client support
 - Middlebox interferences
- Coarse protocol tuning
 - Global parameters (sysctl,...)
 - Very specific socket options



Processing New STAT frames





Proof of Consistency



Plugin Overhead

Intel Xeon E5-2640 v3



Plugin	Mean Goodput -7% good	lput
PQUIC, no plugin	+8% CPU inst	ructions
Monitoring	1037 Mbps	
Multipath 1 path	757 Mbps	
Monitoring + Multipath 1 path	714 Mbps	

JITed eBPF ~2x slower than native code Get/set interface ~5x slower than direct access

Verifying the Termination of Plugins with T2

Plugin	Lines of C Code	# Bytecodes	Proven terminating
Monitoring	500	14	13
VPN	500	11	8
Multipath	2600	32	29
FEC	2500	51	37