

Application of Explicit Measurement Techniques for QUIC Troubleshooting

draft-mdt-quic-explicit-measurements

Alexandre Ferrieux

Igor Lubashev

Giuseppe Fioccola

Marcus Ihlar

Fabio Bulgarella

Mauro Cociglio

Massimo Nilo

Isabelle Hamchaoui

Background

- Spin bit defined in RFC 9000 – can be used to measure roundtrip delays.
 - Desire to measure packet loss
 - Spin bit measurement accuracy degrades with network impairments
- RFC 9506 describes protocol agnostic measurement techniques using explicit bits in packet headers.
- **draft-mdt-quic-explicit-measurements** is an attempt at mapping some of the bits defined in RFC9506 onto QUIC.
- Motivation:
 - Diagnostic measurements, benchmarking, and troubleshooting.
 - Complements lower-layer telemetry by tying measurements to application sessions.

RFC 9506 Recap

		Method	Bits	Unidirectional Observer	Bidirectional Observer	# of Measurements	Impairment resiliency
Delay Bits	S: Spin Bit		1	RTT	Half-RTT x2	Medium	Low
	D: Delay bit		1	RTT	Half-RTT x2	Medium	Medium
		Method	Bits	Unidirectional Observer	Bidirectional Observer	Protocols	Measurement Fidelity
Loss Bits	T round Trip loss bit		1+spin	Round Trip	Round Trip Half-RT x2	All	Rate by sampling
	Q sQuare bit		1	Upstream	Upstream x2	All	Rate over N packets
	L Loss event bit		1	End-to-End	End-to-End x2	with loss detection	Loss shape and rate

Note: The algorithms could also be combined and work independently to improve resiliency and accuracy:

- SD (Spin bit + Delay bit), QL (sSquare + Loss event bits), QR (sSquare + Reflection square bit)

Issue: Real Estate

L bit
Q bit
D bit
Spin bit
...




?

```
1-RTT Packet {  
  Header Form (1) = 0,  
  Fixed Bit (1) = 1,  
  Spin Bit (1),  
  Reserved Bits (2),  
  Key Phase (1),  
  Packet Number Length (2),  
  Destination Connection ID (0..160),  
  Packet Number (8..32),  
  Packet Payload (8..),  
}
```

Issue: coordination / detection

- QUIC v1 header bits are greased.
- QUIC v1 header bits might get used in various ways with different extensions.
- Detecting whether bits are used for measurement is possible but not straight forward.

Solution?



Explicit Flow Measurement Protocol

Modelled on SCONE

- Uses dedicated Long Header Packet with specific version number
- 6 bits of payload in first octet
- QUIC endpoints announce support with transport parameters
- Endpoints insert packet in front of “regular” QUIC v1/v2/vN packets in a UDP datagram
- Connection Ids in measurement packet must be identical to those of inner packets to ensure correct routing

Motivation:

- Explicit measurement layer removes need of reserving space in QUIC Short Header.
- Measurement-enabled flows are easily identified by on-path measurement nodes.
- Lower layer measurement mechanisms are working in similar ways by defining extension headers / encapsulations (e.g., IOAM for IPv6 or MPLS).

Explicit Flow Measurement Packet

```
EFMP Packet {  
  Header Form (1) = 1,  
  Reserved (1),  
  Q Bit (1),  
  L Bit (1),  
  Spin Bit (1),  
  Reserved (3),  
  Version (32) = 0xTBD,  
  Destination Connection ID Length (8),  
  Destination Connection ID (0..2040),  
  Source Connection ID Length (8),  
  Source Connection ID (0..2040),  
}
```

Issues and trade-offs

- Overhead
 - Sent with each packet
 - Packet size at least 8 octets, likely larger, depending on CID length.
- Integrity
 - Measurement packet is completely unprotected.
 - Measurement nodes should not modify packet payload.
 - Endpoints could detect tampering at the cost of extra overhead and complexity.

Issues and trade-offs

- Faked or skewed measurements
 - Endpoints may produce false measurement results by not setting bits according to algorithms.
 - Networks should not trigger policy actions based on measurements.
- Privacy and security
 - Sessions that enable measurements are clearly identifiable.
 - On-path adversaries may use that to track user behaviour.
 - On-path adversaries may want to specifically target measurement traffic to disrupt operators.

Additional Open Issues

- Could this extension remove the spin-bit from QUIC v1/v2 Short Header?
 - Trade-off between overhead for spin-bit only measurement and freeing up bits in Short Header.
- Some methods depend on each packet in a flow being marked, others require less frequent marking.
 - Could specific methods be negotiated?