Explicit Flow Measurements

draft-ietf-ippm-explicit-flow-measurements-00

https://datatracker.ietf.org/doc/draft-ietf-ippm-explicit-flow-measurements/

March 21, 2022, IPPM WG – IETF 113

M. Cociglio - Telecom Italia (TIM)

A. Ferrieux - Orange Lab

G. Fioccola - Huawei Technologies

I. Lubashev - Akamai Technologies

F. Bulgarella - Telecom Italia (TIM)

I. Hamchaoui - Orange Labs

M. Nilo - Telecom Italia (TIM)

R. Sisto - Politecnico di Torino

D. Tikhonov - LiteSpeed Technologies

Explicit Flow Measurements (EFM)

- Explicit Flow Measurement techniques employ few marking bits, inside the header of each packet, for loss and delay measurement (protocol independent and valuable for encrypted header protocols: e.g. QUIC)

- EPM metrics described in this draft:
  - RTT: Delay bit (**D-bit**) (with «the hidden RTT» option: **D^-bit**)
  - Round Trip Packet Loss: Spin bit (**S-bit**) + roundTrip loss bit (**T-bit**)
  - One Way Packet Loss, **2 options**:
    1) sQuare bit (**Q-bit**) + Loss event bit (**L-bit**)
    2) sQuare bit (**Q-bit**) + Reflection square bit (**R-bit**)

IETF Hackathon and Implementations

- Some of the methodologies are already included in ongoing experiments and implementations:
- “QUIC Measurements” project during the last Hackathons
- EFM Implementations in production network reported by the contributors:
  - ***Telecom Italia-TIM*** *Implementation => android mobile phones probe.*
  - ***Ericsson*** *implementation => core network probes.*
  - ***Orange-Akamai*** *implementation => Akamai production CDNs and core network probes.*
  - ***Huawei*** *is working on the topic.*

Universities:

- ***Aachen University*** *implementation: ANRW paper (Packet Loss measurements: L, Q, R, T bits) and Intel Tofino Implementation (paper)*
- ***Politecnico di Torino*** *is developing explicit performance probes.*
- ***Technion*** *(**Israel Institute of Technology**) is developing explicit performance probes.*

The Hidden Bits

In case of strictly privacy concerns it could be introduced the Hidden Delay Bit or the Hidden Spin Bit.

The algorithm of the Delay Bit can be slightly modified to mask the RTT of the connection to an intermediate observer. Even the Spin Bit could be easily modified by implementing the same behaviour.

The idea is to change the Delay Bit working mechanism adding a fixed amount of time to the RTT exposed on wire. In practice, the Client does not reflect immediately the Delay Bit Sample but waits for an Additional Delay before reflecting the marking on a new packet in the opposite direction.

This leads an observer to measure a fake RTT greater than the real one.

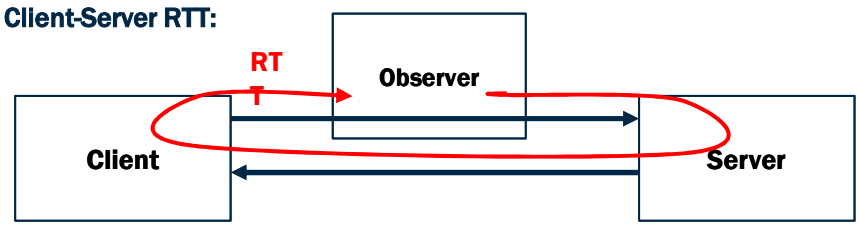

End-to-End Round-Trip Time => RTT = Ts(DbS_2) – Ts(DbS_1) – AD

Ts: Timestamp *DbS: Delay bit Sample* *AD: Additional Delay*

Draft Updates

- **IPPM WG adoption (22/10/2021) :**
- **Updated Draft publication (25/10/2021) :**
  - draft-ietf-ippm-explicit-flow-measurements-00
- **New version next days:**
  - **Ike Kunze (Aachen University) revised all algorythm descriptions because his research group implemented all the measurements.**
    - **Main changes in T_bit description (Round Trip Packet Loss) to clarify token mechanism and measurement period duration.**

“Delay Bits” Summary

	Bits	Unidirectional Observer	Bidirectional Observer	# of Measurements	Impairment resiliency
S: Spin Bit	1	RTT	x2 Half-RTT	Very High	Low
D: Delay bit	1	RTT	x2 Half-RTT	Medium °	High
D^: Hidden Delay bit	1	RTT^	x2 Left Half-RTT^ Right Half-RTT	High ~	High
SD: Spin bit + Delay bit *	2	RTT	x2 Half-RTT	Very High	High

° It depends on the “application delay” threshold (e.g. E=1 ms.), causing DbS discarded, and on DbS losses. But many of these missing measurements are “errored” measurements.

~ The “application delay” threshold (e.g. E=1 ms.) is only on the Server (see previous note).

* Both algorithms work independently; an observer could use approximate spin bit measures when delay bit ones aren't available.

X2 Same metric for both directions.

^ Masked metric (real value can be calculated only by those who know the Additional Delay).

I E T F®

“Loss Bits” Summary

Method	Bits	Unidirectional Observer	Bidirectional Observer	Proto	Measurement Fidelity	Measurement Delay
T round Trip loss bit	1+spin	Round Trip	Round Trip Half-RT x2	*	Rate by sampling $\frac{1}{3}$ to $\frac{1}{3*ppa}$ packets over 2 RTT	~6 RTT
Q sQuare bit	1	Upstream	Upstream x2	*	Rate over N packets (e.g. N=64)	N packets (e.g. B-64)
L Loss event bit	1	End-to-End	End-to-End x2	#	Loss shape and rate	Min: RTT Max: RTO
QL sQuare + Loss event bits	2	Upstream Downstream End-to-End	Upstream x2 Downstream x2 End-to-End x2	#	→ see Q → see Q|L → see L	→ see Q → see L → see L
QR sQuare + Reflection square bit	2	Upstream “3/4 RT” Opp. Dir. E2E	Upstream x2 “3/4 RT” x2 End-to-End x2 Downstream x2 Half-RT x2	*	Rate over $N * ppa$ packets (e.g. N=64)	Upstream: see Q Others: $N * ppa$ pkts (e.g. N=64)

*	All protocols	x2	Metric in both directions	ppa	Packets-per-Ack
#	Protocols with loss detection (w/ or w/o pure ACK loss detection)			Q|L	See Q if Upstream loss is significant; L otherwise

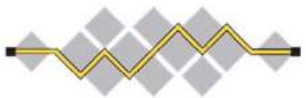

2 or 3bit EFM

If there are only 2 bits for EFM:

- Option a:
  - RTT (S-bit)
  - RT Packet Loss (T-bit)
- Option b:
  - RTT (D-bit or D^-bit)
  - OneWay P.Loss (Q-bit)

If there are 3 bits for EFM (e.g. QUIC):

- Option c:
  - RTT (S-bit or S^-bit)
  - OneWay P. Loss (Q-bit + L-bit)
- Option d:
  - RTT (D-bit or D^-bit)
  - OneWay P.Loss (Q-bit + L-bit)
- Option e:
  - RTT (S-bit or S^-bit)
  - OneWay P. Loss (Q-bit + R-bit)
- Option f:
  - RTT (D-bit or D^-bit)
  - OneWay P.Loss (Q-bit + R-bit)

Conclusions

- **Explicit Flow Measurements are gaining interest for encrypted transport protocols:**
  - implementation at IETF Hackathon;
  - On field implementations
  - WG adoption
  - Academic analysis
- **Sibling Draft in IPPM WG**
  - https://datatracker.ietf.org/doc/draft-cnbf-ippm-user-devices-explicit-monitoring/
- **Sibling Drafts in other WGs (COAP, QUIC, TCP, ...)**
  - https://datatracker.ietf.org/doc/draft-fz-core-coap-pm/
  - **...**
- **Next step WG Last Call**

Welcome questions and comments!

Thank you

User Devices Explicit Monitoring

draft-cnbf-ippm-user-devices-explicit-monitoring-03

https://datatracker.ietf.org/doc/draft-cnbf-ippm-user-devices-explicit-monitoring/

M. Cociglio - Telecom Italia (TIM)

F. Bulgarella - Telecom Italia (TIM)

M. Nilo - Telecom Italia (TIM)

G. Fioccola - Huawei Technologies

Proposal: EFM Probes on user devices

- The draft proposes to put the Explicit Flow Measurements probe also on the user device (e.g. mobile phones, PCs).
- "User device EFM rules":
  1. **The device owner decides whether to mark his traffic.**
  2. **The device owner decides whether to share his performance data.**
- Strenghts:
  1. **Scalability**. On the user device there are few connections to monitor.
  2. **More precise measurements**. Client application delay can be measured.
  3. **Both directions monitoring.**
  4. **Network monitoring equipment savings**. Network probes can monitor only impaired connections through "**user device and network probes coordination**". *It's possible to set alarm thresholds on the user device (and to signal to network probes to monitor only the sessions with impairments, in order to segment the performance measurements and to locate the faults). In this case network probes, also embedded into network nodes, need to monitor only a limited number of connections.*

Explicit Flow Measurements

Protocol implementations: the bit choice (a first proposal)

M. Cociglio - Telecom Italia (TIM)

F. Bulgarella - Telecom Italia (TIM)

M. Nilo - Telecom Italia (TIM)

IETF

2bit EFM, possible solutions

If there are only 2 bits for EFM (COAP)

Preferred option:

- RTT (S-bit)
- RT Packet Loss (Q-bit)

Backup option:

- RTT (D-bit)
- RT Packet Loss (L-bit or T-bit)

3bit EFM, suggested solutions

QUIC

Preferred option:

- RTT (D^-bit)
- OneWay P. Loss (Q-bit + L-bit)

Backup option:

- RTT (S^-bit)
- OneWay P. Loss (Q-bit + R-bit)

TCP:

- RTT (D-bit)
- OneWay P. Loss (Q-bit + R-bit)

D^-bit: Hidden Delay bit
S^-bit: Hidden Spin bit

Loss bits strengths

L-bit: simpler implementation, less measurement delay

R-bit: detect losses also for all ack TCP packets, protocol independent