

# An Upgrade to Benchmarking Methodology for Network Interconnect Devices

**Gábor Lencse**

Budapest University of Technology and Economics

**Keiichi Shima**

IIJ Innovation Institute

IETF 108, BMWG, July 27, 2020.

# What can be updated?

- Recommendation to backport the novelties of RFC 8219 to RFC 2544
- Improved throughput and frame loss rate measurement procedures using individual frame timeout
- Requirement of statistically relevant number of tests
- An optional non-zero frame loss acceptance criterion for the throughput measurement procedure

# Novelties of RFC 8219

- Compared to RFC 2544 and RFC 5180
  - New measurement procedures:
    - PDV: Packet Delay Variation
    - IPDV: Inter Packet Delay variation
  - Different summarizing functions
    - RFC 2544: average (a single number: oversimplification)
    - RFC 8219: median plus 1<sup>st</sup> and 99<sup>th</sup> percentiles
  - Higher statistical reliability
    - Requirement for at least 20 tests
  - Redefined Latency measurement procedure
    - At least 500 timestamps instead of a single one

# Why not backporting?

- RFC 8219 is for IPv6 transition technologies.
- The new or redefined measurement procedures can be applied to any network interconnect devices, too.
- Let us update RFC 2544 to do so!
- There is a free software Tester that supports the new PDV, IPDV and the redefined Latency tests
  - **siitperf** (RFC 8219 compliant SIIT tester) can be configured to benchmark IPv4 or IPv6 routers, too.
  - <https://github.com/lencsegabor/siitperf>

# Improved throughput measurements

- RFC 2544 throughput measurements procedure
  - counts the sent and received frames
  - "timeout" is 62s/2s for the first/last frame
  - It is usually OK for hardware forwarding devices
  - But may be a problem, if software solutions introduce (selective) high latencies

- As shown by our experimental results: 100ms delay to 1% of the test frames caused more than 50% decrease in the throughput of HTTP download

G. Lencse, K. Shima, and A. Kovács, "Gaming with the Throughput and the Latency Benchmarking Measurement Procedures of RFC 2544", *International Journal of Advances in Telecommunications, Electrotechnics, Signals and Systems*, vol. 9, no. 2, pp. 10-17, 2020, DOI: 10.11601/ijates.v9i2.288

# Improved throughput measurements

- Recommended solution
  - An advanced throughput measurement procedure that checks the timeout time (e.g. 10ms) for every single test frame
    - We have demonstrated its feasibility with **siitperf**  
G. Lencse, "Design and Implementation of a Software Tester for Benchmarking Stateless NAT64 Gateways", accepted for *IEICE Transactions on Communications*, available: <http://www.hit.bme.hu/~lencse/publications/IEICE-2020-siitperf-revised.pdf>
    - The value of the "frame timeout" is subject to research
    - Question: Does it make a significant difference?
      - <https://mailarchive.ietf.org/arch/msg/bmwg/50qoL0gxTEKGU6CkUwPIf8FO-hc/>
- The same can be applied to frame loss rate tests, too.

# Requirement of statistically relevant number of tests

- RFC 8219 mentions at four different places that the tests must be repeated at least 20 times:
  - latency (Section 7.2)
  - packet delay variation (Section 7.3.1)
  - inter packet delay variation (Section 7.3.2)
  - DNS64 performance (Section 9.2).
- On the one hand, a similar guideline could be nice for the throughput test, too.
- On the other hand, the binary search is time consuming: perhaps 20 repetitions is too many.

# Requirement of statistically relevant number of tests

- Our recommendation:
  - To develop an algorithm that checks the statistical properties of the results of the tests
    - It may stop before 20 repetitions, if the results are consistent,
    - it may require more than 20 repetitions, if the results are scattered.



# An optional non-zero frame loss acceptance criterion for the throughput measurement procedure

- Arguments:
  - Packet forwarding is often implemented in software. It is not feasible to require 0% frame loss.
  - Applications usually tolerate some low frame loss rates (e.g. 0.01%)
  - Commercial Testers usually allow to specify "Loss Tolerance"
  - It is better to allow such measurements and to require stating the applied loss tolerance rate.

# Thank you for listening!

- The Internet Draft is available:

<https://tools.ietf.org/html/draft-lencse-bmwg-rfc2544-bis-00>

- All comments are welcome!
- Our question:
  - Do you consider any of our recommendations useful?