Network Working Group
Internet-Draft
Updates: 2544 (if approved)
Intended status: Informational
Expires: November 20, 2020

A. Morton
AT&T Labs

May 19, 2020

Updates for the Back-to-back Frame Benchmark in RFC 2544

draft-ietf-bmwg-b2b-frame-02

Abstract

Fundamental Benchmarking Methodologies for Network Interconnect Devices of interest to the IETF are defined in RFC 2544. This memo updates the procedures of the test to measure the Back-to-back frames Benchmark of RFC 2544, based on further experience.

This memo updates Section 26.4 of RFC 2544.

Addressing Vratko's late November 2019 and May 2020 Comments:

- 02 Proposed resolutions
- Buffer size expressed as time in the original tests
- Buffer Time defined
- Buffer Filling Rate defin.
- Average Repeated Tests
- Revised Time Correction Factor to calculate Time

- Support expressed in May
- Keep in mind:
  - This is one of several updates to RFC2544
- Previous updates include:
  - IPv6 address space
  - Reset & Restoration
  - MPLS
  - State of the art Latency

When the VSPERF CI results were examined [VSPERF-b2b], several aspects of the results were considered notable:

1. Back-to-back Frame Benchmark was very consistent for some fixed frame sizes, and somewhat variable for others.

2. The number of Back-to-back Frames with zero loss reported for large frame sizes was unexpectedly long (translating to 30 seconds of buffer time), and no explanation or measurement limit condition was indicated. It's important that the buffering time was used in this analysis. The referenced testing [VSPERF-b2b] and calculations produced buffer extents of 30 seconds for some frame sizes, and clearly wrong in practice. On the other hand, a result expressed only as a large number of Back-to-back Frames does not permit such an easy comparison with reality.

3. Calculation of the extent of buffer time in the DUT helped to explain the results observed with all frame sizes (for example, some frame sizes cannot exceed the frame header processing rate of the DUT and therefore no buffering occurs, therefore the results depended on the test equipment and not the DUT).

4. It was found that the actual extent of buffer time in the DUT could be estimated using results to measure the longest burst in frames without loss and results from the Throughput tests conducted according to Section 26.1 of [RFC2544]. It is apparent that the DUT's frame processing rate empties the buffer during a trial and tends to increase the "implied" buffer size estimate (measured according to Section 26.4 of [RFC2544] because many frames have left the buffer when the burst of frames ends). A calculation using the Throughput measurement can reveal a "corrected" buffer size estimate.

In reality, there are many buffers and packet header processing steps in a typical DUT. The simplified model used in these calculations for the DUT includes a packet header processing function with limited rate of operation, as shown below:

```
                    |------------ DUT --------|
Generator -> Ingress -> Buffer -> HeaderProc -> Egress -> Receiver
```

So, in the back2back frame testing:

1. The Ingress burst arrives at Max Theoretical Frame Rate, and initially the frames are buffered.

2. The packet header processing function (HeaderProc) operates at the "Measured Throughput" (Section 26.1 of [RFC2544]), removing frames from the buffer (this is the best approximation we have).

3. Frames that have been processed are clearly not in the buffer, so the Corrected DUT buffer time equation (Section 5.4) estimates and removes the frames that the DUT forwarded on Egress during the burst. We define buffer time as the number of Frames occupying the buffer divided by the Maximum Theoretical Frame Rate (on egress) for the Frame size under test.

4. A helpful concept is the buffer filling rate, which is the difference between the Max Theoretical Frame Rate (ingress) and the Measured Throughput (HeaderProc on egress). If the actual buffer size in frames was known, the time to fill the buffer during a measurement can be calculated using the filling rate as a check on measurements. However, the Buffer in the model represents many buffers of different sizes in the DUT data path.

The original Section 26.4 of [RFC2544] definition is stated below:

> The Back-to-back Frame value is the longest burst of frames that the DUT can successfully process and buffer without frame loss, as determined from the series of trials.

5.3. Test Repetition and Benchmark

On this topic, Section 26.4 of [RFC2544] requires:

> The trial length MUST be at least 2 seconds and SHOULD be repeated at least 50 times with the average of the recorded values being reported.

Therefore, the Benchmark for Back-to-back Frames is the average of burst length values over repeated tests to determine the longest burst of frames that the DUT can successfully process and buffer without frame loss. Each of the repeated tests completes an independent search process.

In this update, the test MUST be repeated N times (the number of repetitions is now a variable that must be reported),for each frame size in the subset list, and each Back-to-back Frame value made available for further processing (below).

The next step is to apply a correction factor that accounts for the DUT's frame forwarding operation during the test (assuming the simple model of the DUT composed of a buffer and a forwarding function, described in Section 3).

Corrected DUT Buffer Time =

$$= \text{Implied DUT Buffer Time} - \left(\text{Implied DUT Buffer Time} * \frac{\text{Measured Throughput}}{\text{Max Theoretical Frame Rate}} \right)$$

where:

1. The "Measured Throughput" is the [RFC2544] Throughput Benchmark for the frame size tested, as augmented by methods including the Binary Search with Loss Verification aglorithm in [TST009] where applicable, and MUST be expressed in Frames per second in this equation.

Next Steps

- Trigger any concluding reviews with a WG Last Call!