Short Performance Evaluation Report

The Test Systems

My virtual machine (VM):

Intel® Xeon® CPU E5-2680, 16 Cores, 32GB RAM

My windows subsystem for Linux (WSL)

Intel® Core ™ i7-1065G7, 8 Cores, 32GB RAM

The test

We use a local dockerized tiny topology. We have a set of 1024 packets where all of them are identical except for the flow ID where each packet has an increasing flow ID. We use tcpreplay to send those 1024 packets in a loop to the border router as fast as possible and query the metrics endpoint of the border router every second. This is repeated for 30 seconds. This test is performed for a payload size of 100 bytes, 500 bytes and 1000 bytes on both test systems and compared to the performance of the old border router with the same test data.

Observations

The performance was optimized between the old border router and the new border router. The current bottleneck, as can be seen in the pprof files, lies in the receiver and forwarder, mostly in the ReadBatch() and WriteBatch() calls.

The results

The packet sizes are measured for the entire frame. 222 bytes, 622 bytes and 1122 bytes respectively.

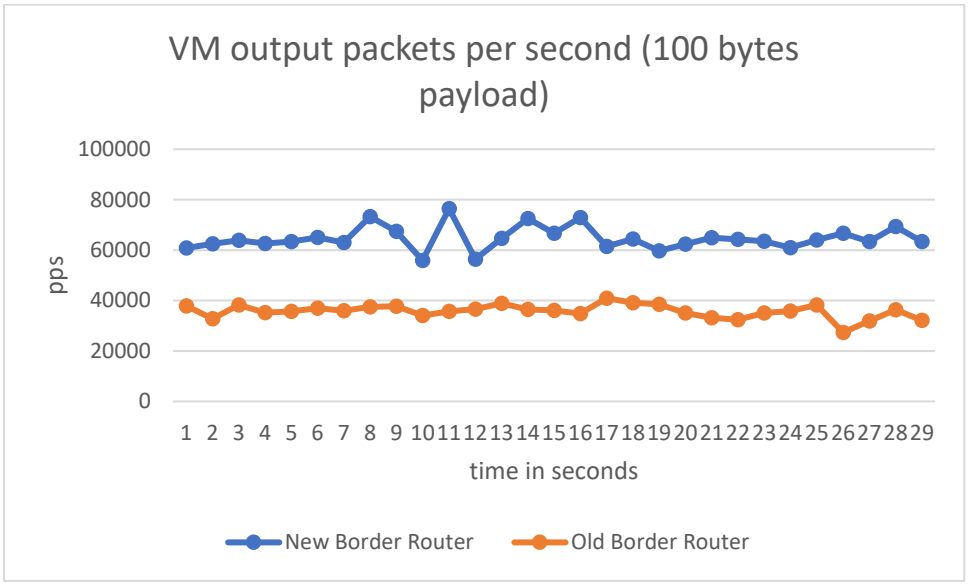

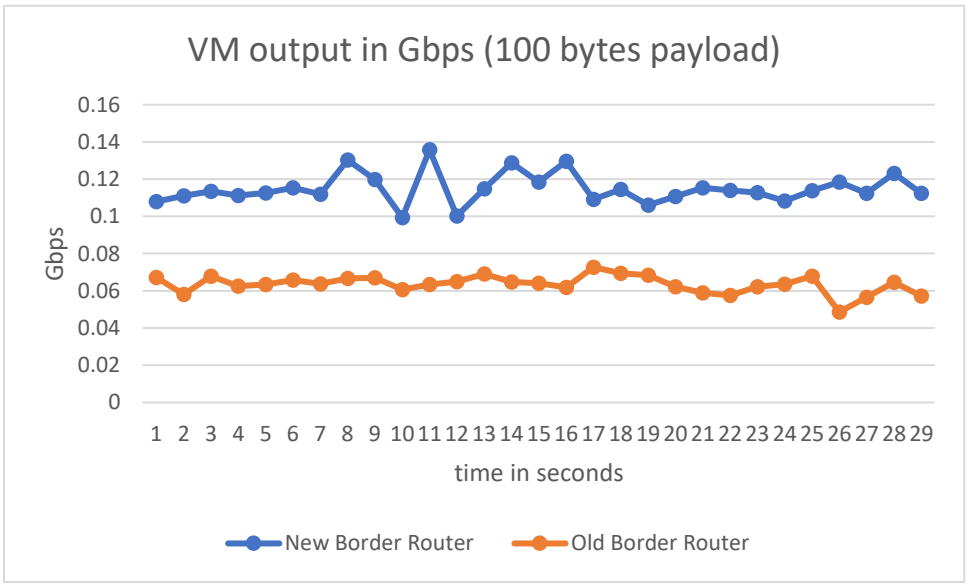

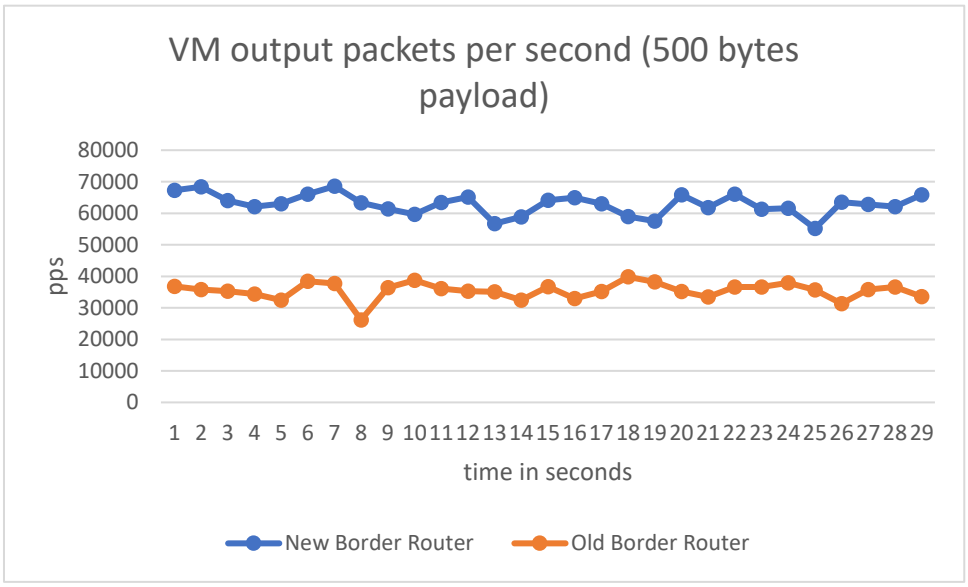
VM output packets per second (500 bytes payload)
pps
80000
70000
60000
50000
40000
30000
20000
10000
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

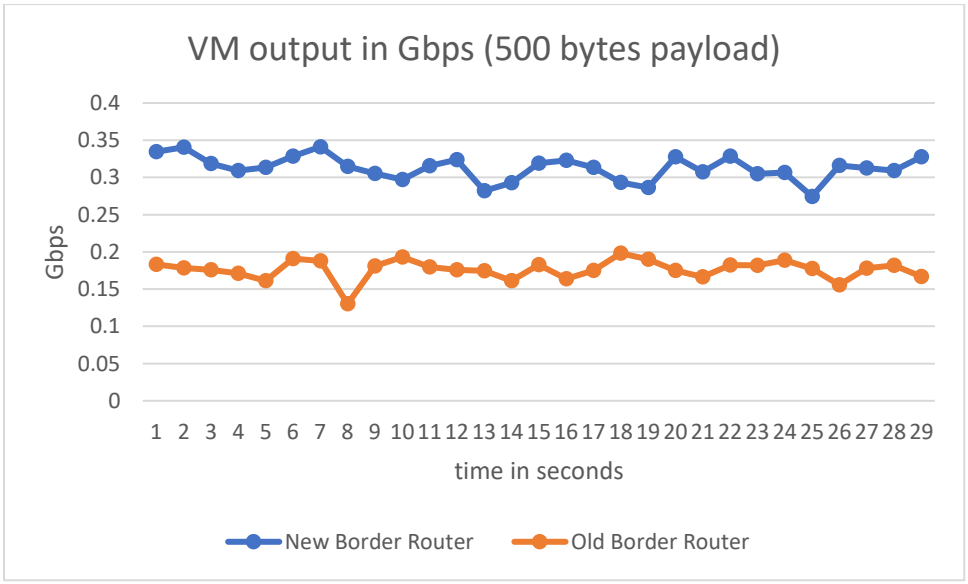
VM output in Gbps (500 bytes payload)
Gbps
0.4
0.35
0.3
0.25
0.2
0.15
0.1
0.05
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

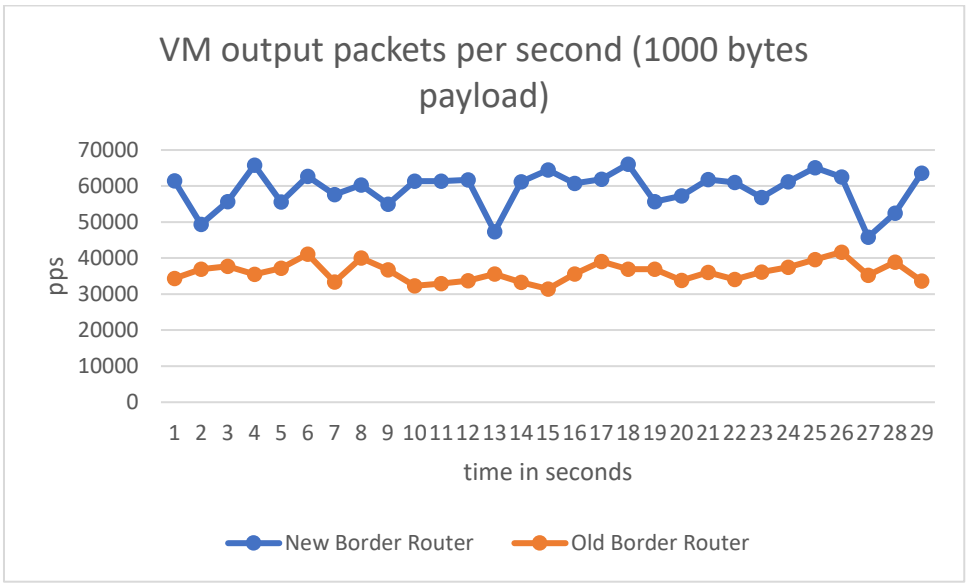
VM output packets per second (1000 bytes payload)
pps
70000
60000
50000
40000
30000
20000
10000
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

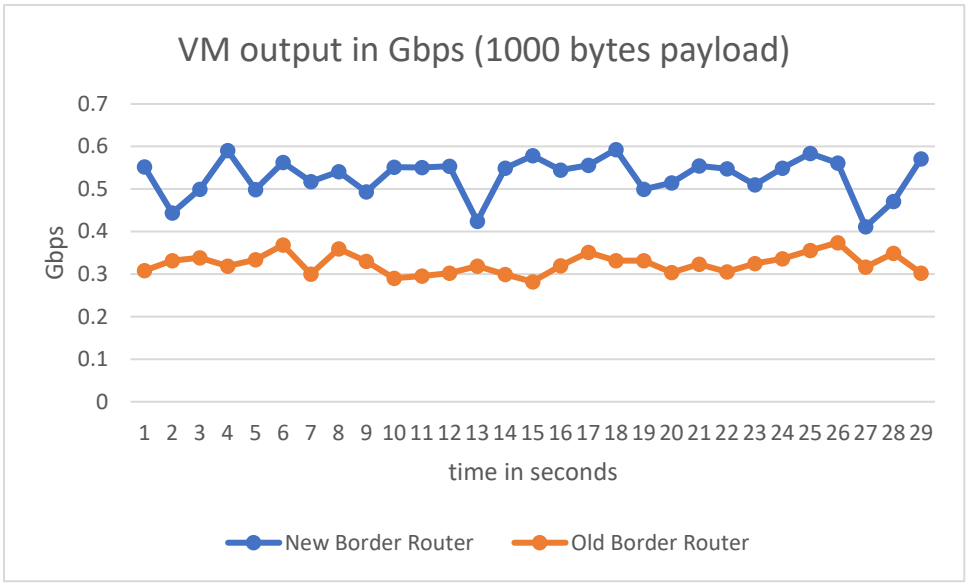
VM output in Gbps (1000 bytes payload)
Gbps
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

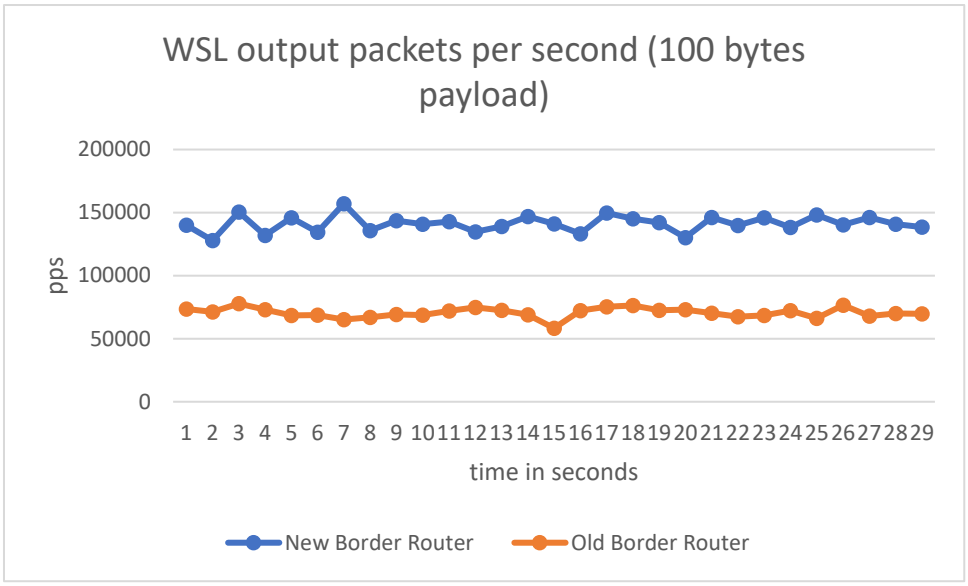
WSL output packets per second (100 bytes payload)
200000
150000
100000
50000
0
pps
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

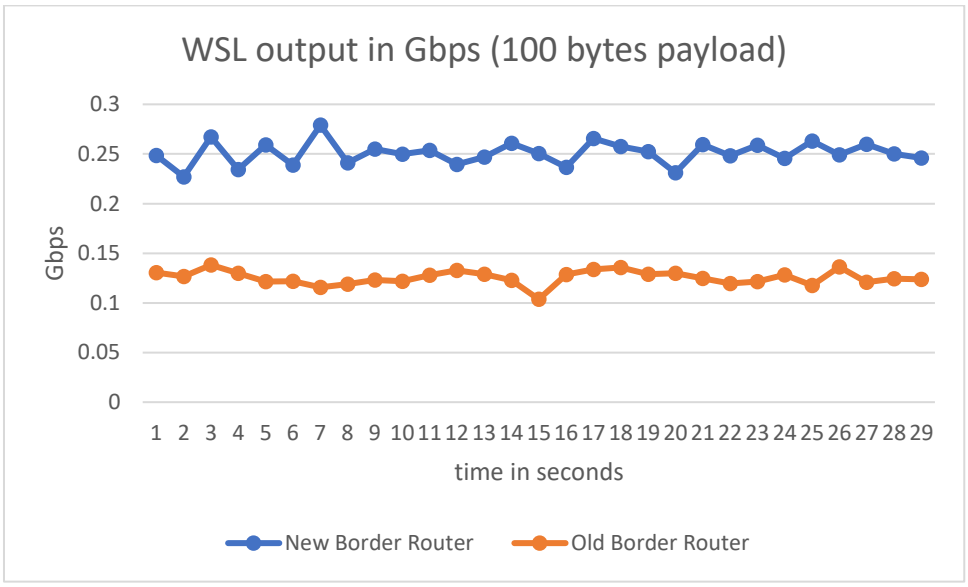
WSL output in Gbps (100 bytes payload)
0.3
0.25
0.2
0.15
0.1
0.05
0
Gbps
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

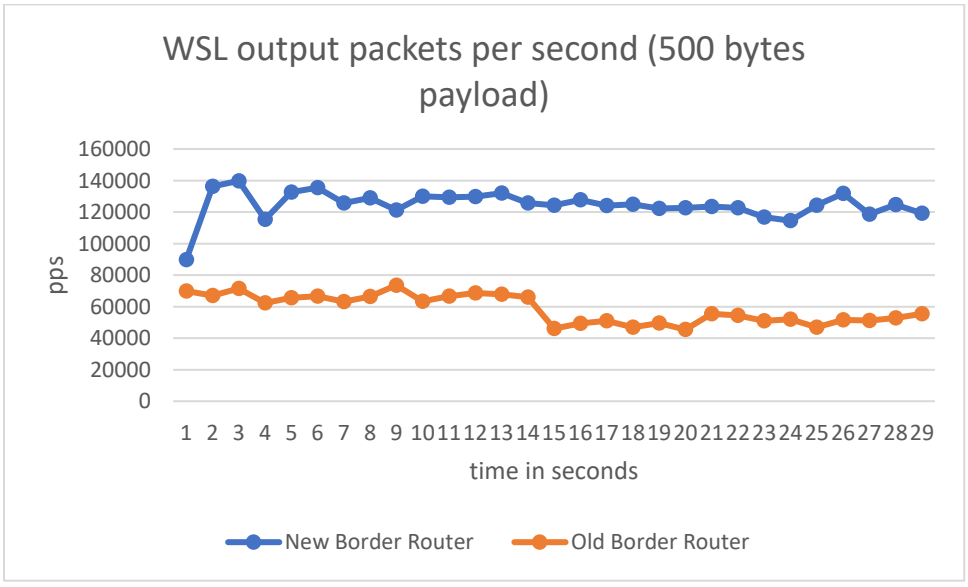
WSL output packets per second (500 bytes payload)
160000
140000
120000
100000
80000
60000
40000
20000
0
pps
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

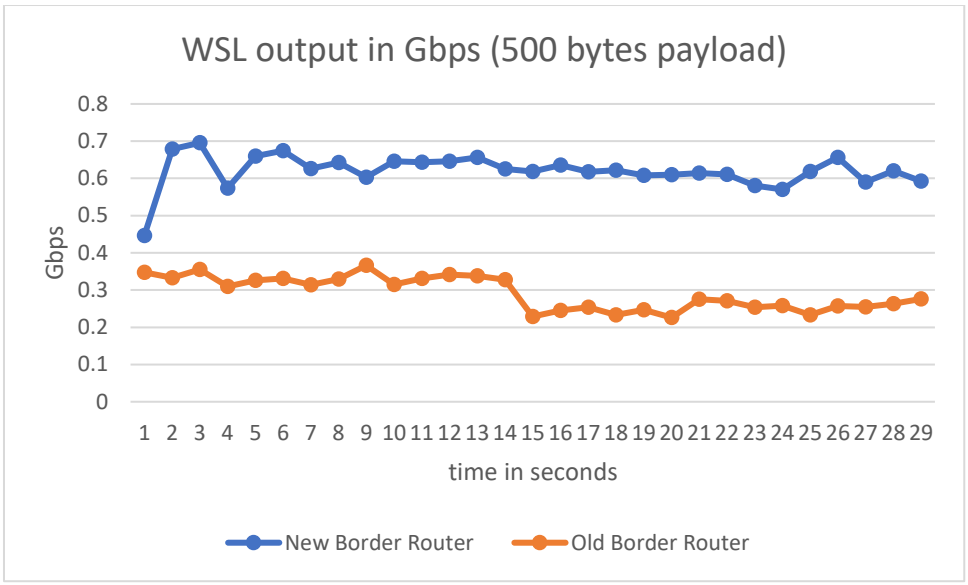
WSL output in Gbps (500 bytes payload)
Gbps
0.8
0.7
0.6
0.5
0.4
0.3
0.2
0.1
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

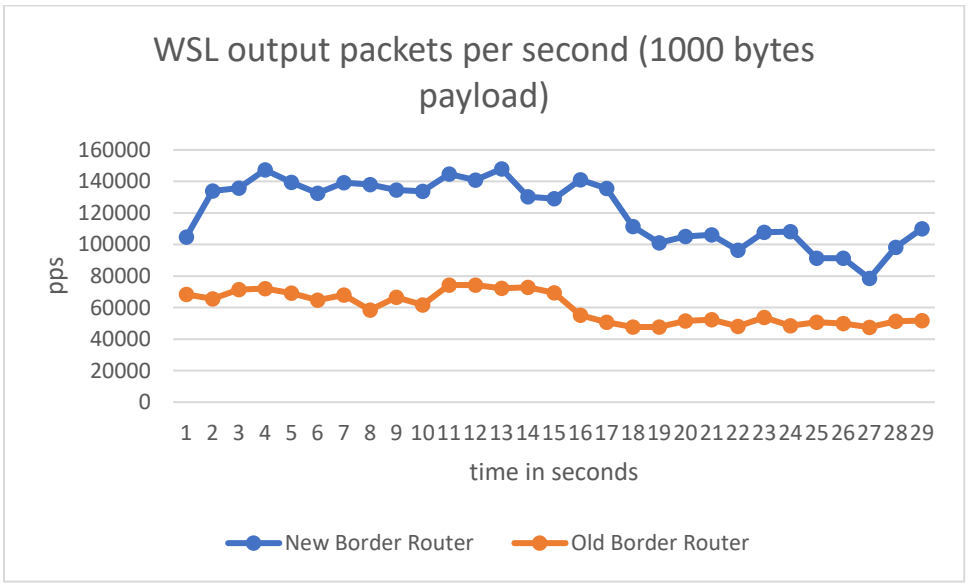
WSL output packets per second (1000 bytes payload)
pps
160000
140000
120000
100000
80000
60000
40000
20000
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router

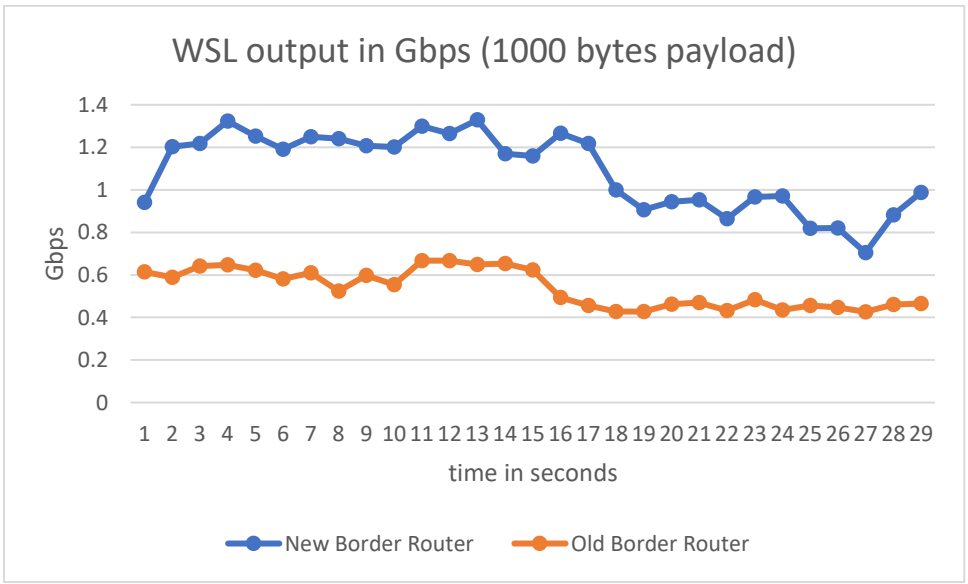
WSL output in Gbps (1000 bytes payload)
Gbps
1.4
1.2
1
0.8
0.6
0.4
0.2
0
1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29
time in seconds
New Border Router
Old Border Router