

Evaluation of Broadcom NetXtreme 25Gb Ethernet Adapters

Broadcom demonstrates the industry's most deterministic latency – the critical element to providing consistent, scalable performance for enterprise and cloud environments.



Executive Summary

As customers of all sizes begin to adopt 25Gb Ethernet (25GbE) technology, enterprises and cloud data centers need to consider the scalability and consistency of performance of the Ethernet adapters, also known as network interface cards (NICs).

In particular, scalable data centers, where hundreds to thousands of NICs are deployed, need to consider the entire latency distribution including tail latency – worst-case latency toward the latter end of the latency distribution.

Tail latency matters more than minimum or average latency, because the tail latency experienced by several NICs at any given moment can cause unexpected and disruptive application delays, slowing down the overall network performance. Low and predictable latency is critical to high-performance scalable data centers. This dynamic is highlighted in the paper *The Tail at Scale* by Dean and Barroso from Google.¹

Broadcom commissioned Demartek to analyze the performance of the Broadcom® NetXtreme® 25GbE NIC in terms of performance consistency, comparing these results to the equivalent Mellanox® ConnectX®-4 25GbE NIC.

Latency and Throughput Metrics in Real-World Environments

In this paper, Demartek measures two of the most common NIC performance metrics: throughput and latency.

Most traditional benchmark tests measure performance under little to no traffic. However, this is unrealistic and generally not helpful in measuring real-world performance.

In the testing performed, we focused on the following:

- > Realistic network traffic scenarios
 - > Latency measured with loaded traffic
 - > For RoCEv2 tests, congestion was injected to mimic a realistic network scenario
- > Latency that matters
 - > Average latency
 - > Latency variability

Key Findings

- > Broadcom's 25GbE adapter shows lower average latency.
- > Broadcom's 25GbE adapter provides higher lossless frame rates than the competitive adapter.
- > Broadcom's 25GbE adapter has very small latency variability, while the competitive adapter has very high long-tail latency.
- > In RoCEv2 tests, Broadcom's 25GbE adapter shows higher port throughput and more consistent latency than the competitive adapter under network congestion.

¹ <https://research.google.com/pubs/pub40801.html>

Evaluation of Broadcom NetXtreme 25Gb Ethernet Adapters

RFC2544 – L2 Performance Tests

RFC2544 tests are commonly used to evaluate Ethernet adapters for throughput and latency in transmit and receive tests. For these tests, dual-port 25GbE adapters were tested using Ixia test equipment and the Ixia IxNetwork tools. IPv4 traffic was transmitted into one port of the adapter and returned via the second port.

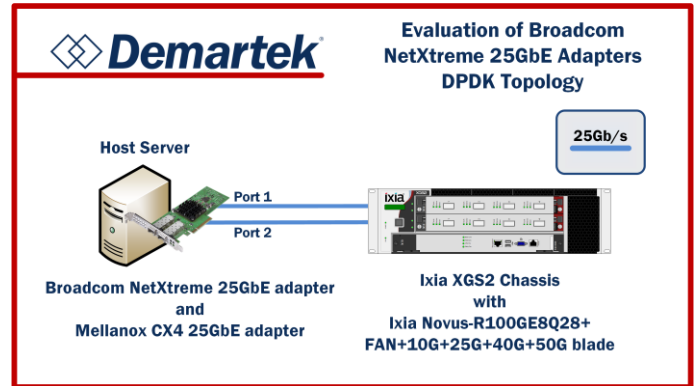

Throughput tests determine the maximum rate at which no offered frames are dropped.

Performance Results – RFC2544 (L2 Tests)

These tests handle packet routing between the test hardware and the adapter installed in the host server, focusing on small packets and utilization the adapter, driver and IP portion of the TCP/IP stack.

Long-tail Latency

The Broadcom adapter shows only 35K packets (out of 100M – approximately three seconds of line rate operation) outside 20 μsec and all packets within 40 μsec, while the competition shows 1.3M packets outside 10 μsec and 145K packets still not received after 120 μsec.



Summary and Conclusion

For large-scale installations, it is important to have consistent performance across all of the adapters so that applications provide predictable performance across hundreds or thousands of nodes. This consistency applies to frames per second, throughput and latency.

With a large number of network adapters deployed in an enterprise or cloud data center, the probability that one or more of these adapters is experiencing its maximum latency at any given moment is high. Wildly fluctuating latencies can lead to, at best, uncertainty for application performance and, at worst, application failures.

In the tests conducted for this report, the Broadcom NetXtreme 25GbE network adapter provided not only lower average latency but a much tighter set of latency results from minimum to maximum latency than the competitive adapter. Moreover, Broadcom's maximum latency under heavy load was up to five times lower than the competitive adapters.

The most current version of this report is available at www.demartek.com/Broadcom-25G on the Demartek website.

Demartek is a registered trademark of Demartek, LLC.

All other trademarks are the property of their respective owners.