

Need for Speed: Comparing FDR and EDR InfiniBand (Part 2)

By Olumide Olusanya and Munira Hussain

This is the second part of this blog series. In the [first part](#), we shared OSU Micro-Benchmarks (latency and bandwidth) and HPL performance between FDR and EDR Infiniband. In this part, we will further compare performance using additional real-world applications such as ANSYS Fluent, WRF, and NAS Parallel Benchmarks. For my cluster configuration, please refer to [part 1](#).

Results

ANSYS Fluent

Fluent is a Computational Fluid Dynamics (CFD) application used for engineering design and analysis. It can be used to simulate the flow of fluids, with heat transfer, turbulence and other phenomena, involved in various transportation, industrial and manufacturing processes.

For this test we ran Eddy_417k which is one of the problem sets from ANSYS Fluent Benchmark suits. It is a reaction flow case based on the eddy dissipation model. In addition, it has around 417,000 hexahedral cells and is a small dataset with a high communication overhead.

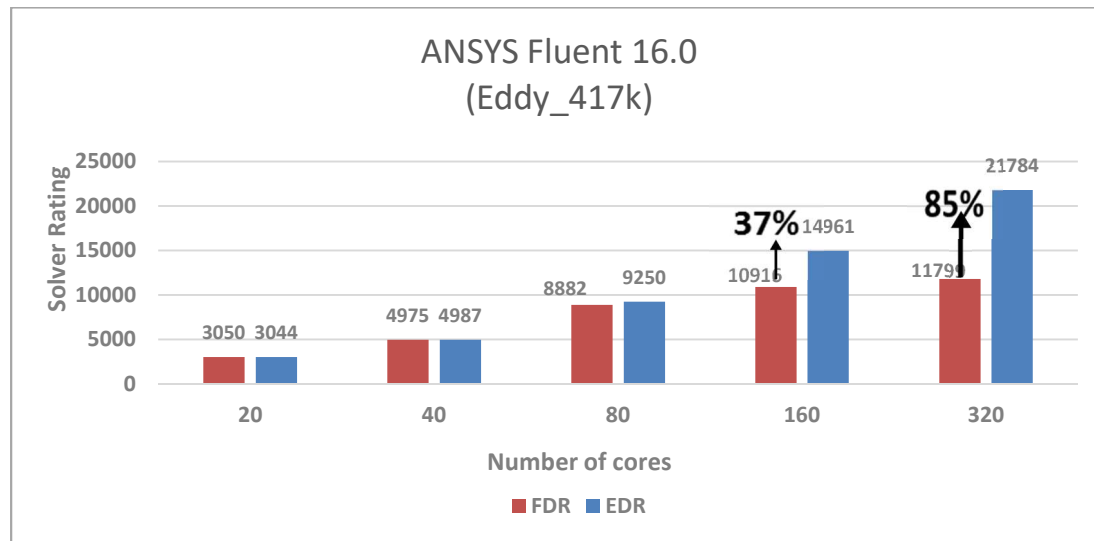


Figure 1 - ANSYS Fluent 16.0 (Eddy_417k)

From **Figure 1** above, EDR shows a wide performance advantage over FDR as the number of cores increase to 80. We continue to see an even wider difference as the cluster scales. While FDR's performance seems to gradually taper off after 80 cores, EDR's performance continues to scale as the number of cores increase and performs 85% better than FDR on 320 cores (16 nodes).

WRF (Weather Research and Forecasting)

WRF is a modelling system for weather prediction. It is widely used in atmospheric and operational forecasting research. It contains two dynamic cores, a data assimilation system, and a software architecture that allows for parallel computation and system extensibility. For this test, we are going to study the performance of a medium size case, Conus 12km.

Conus 12km is a resolution case over the Continental US domain. The benchmark is run for 3 hours after which we take the average of the time per time step.

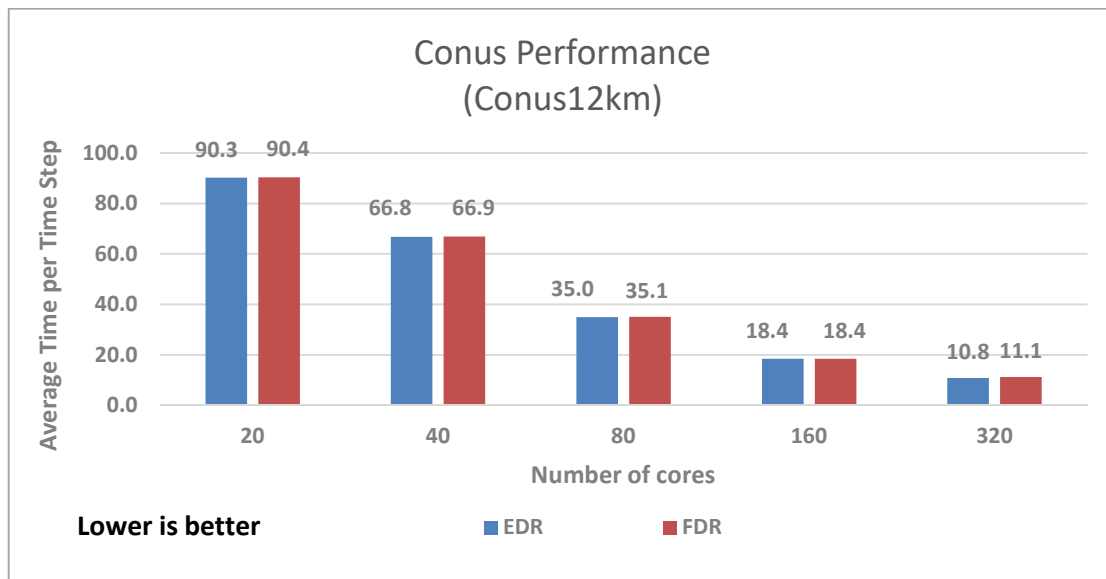


Figure 2 - WRF (Conus12km)

Figure 2 shows both EDR and FDR scaling almost linearly and also performing almost equally until the cluster scales to 320 cores when EDR performs better than FDR by 2.8%. This performance difference, which may seem little, is significantly higher than my highest run to run variation of 0.005% between three successive EDR and FDR 320-core tests.

HPC Advisory Council's result [here](#) shows a similar trend with the same benchmark. From their result, we can see that the performances are neck and neck until the 8 and 16-node run where we see a small performance gap. Then the gap widens even more in the 32-node run and EDR posts a 28% better performance than FDR. Both results show that we could see an even higher performance advantage with EDR as we scale beyond 320 cores.

NAS Parallel Benchmarks

NPB contains a suite of benchmarks developed by NASA Advanced Supercomputing Division. The benchmarks are developed to test the performance of highly parallel supercomputers which all mimic large-scale and commonly used computational fluid dynamics applications in their computation and

data movement. For my test, we ran four of these benchmarks: CG, MG, FT, and IS. In the figures below, the performance difference is in an oval right above the corresponding run.

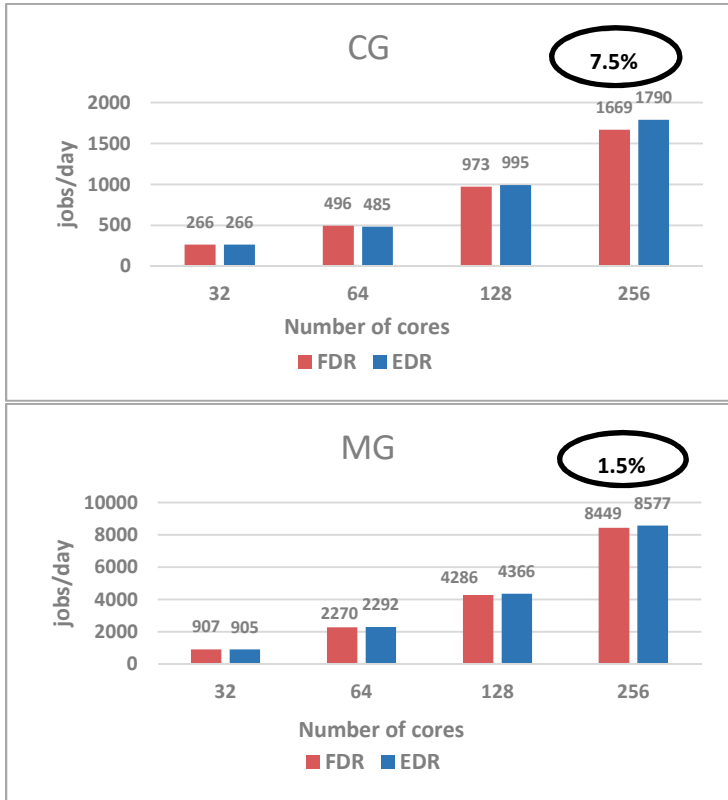


Figure 3 - CG

Figure 4 - MG

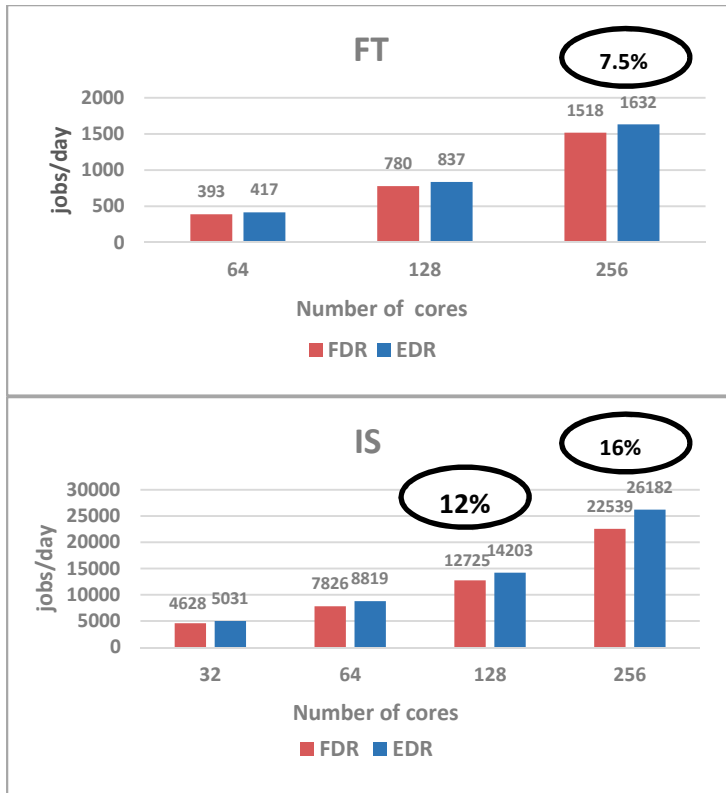


Figure 5 - FT
Figure 6 - IS

CG is a benchmark which computes an approximation of the smallest eigenvalue of a large, sparse, symmetric positive-definite matrix using a conjugate gradient method. It also tests irregular long distance communication between cores. From **Figure 3** above, EDR shows a 7.5% performance advantage with 256 cores.

MG solves a 3-D Poisson Partial Differential Equation. The problem in this benchmark is simplified as it has constant instead of variable coefficients to better mimic real applications. In addition to this, it tests short and long distance communication between cores. Unlike CG, the communication patterns are highly structured. From **Figure 4**, EDR performs better than FDR by 1.5% on our 256-core cluster.

FT is a 3-D partial differential equation solution using FFTs. It tests the long-distance communication performance as well and shows a 7.5% performance gain using EDR on 256 cores as seen in **Figure 5** above.

IS, a large integer sort application, shows a high 16% performance difference between EDR and FDR on 256 cores. This application not only tests the integer computation speed, but also the communication performance between cores. From **Figure 6**, we can see a 12% EDR advantage with 128 cores which increases to 16% on 256 cores.

Conclusion

In both blogs, we have shown several micro-benchmark and real-world application results to compare FDR with EDR Infiniband. From these results, EDR has shown a higher performance and better scaling than FDR on our 16-node Dell PowerEdge C6320 cluster. Also, some applications have shown a wider performance margin between these interconnects than other applications. This is because of the nature of the applications being tested; communication intensive applications will definitely perform and scale better with a faster network when compared with compute-intensive applications. Furthermore, because of our cluster size, we were only able to test the scalability of the applications on 16 servers (320 cores). In the future, we plan on running these tests again on a larger cluster to further test the performance difference between EDR and FDR.