



The Effect of HDR InfiniBand and In-Network Computing on OpenFOAM Simulations

Ophir Maor, Yong Qin, Gerardo Cisneros-Stoianowski, Gilad Shainer
HPC Advisory council

High-performance computing (HPC) technologies are used in the engineering and automotive design and manufacturing industry. One of the applications is the computer-aided engineering (CAE), from component-level design to full analyses such as: crash simulations, structure integrity, thermal management, climate control, modeling, acoustics, and much more. HPC helps drive faster time-to-market, realizing significant cost reductions over laboratory testing and tremendous flexibility. HPC's strength and efficiency depend on the ability to achieve sustained top performance by driving the CPU performance toward its limits. The motivation for high-performance computing has long been its tremendous cost savings and product improvements; the cost of a high-performance compute cluster can be just a fraction of the price of a single crash test for example, and the same cluster can serve as the platform for every test simulation going forward.

The recent trends in cluster environments, such as multi-core CPUs, GPUs, and advanced high speed, low latency interconnect with offloading capabilities, are changing the dynamics of cluster-based simulations. Software applications are being reshaped for higher degrees of parallelism and multithreading, and hardware is being reconfigured to solve new emerging bottlenecks to maintain high scalability and efficiency. Applications like OpenFOAM others are widely used and provide better flexibility, scalability, and efficiency for such simulations, allowing for larger problem sizes and speeding up time to results.

HPC Applications relies on Message Passing Interface (MPI), the de-facto messaging library for high performance clusters that is used for node-to-node inter-process communication (IPC). MPI relies on a fast, unified server and storage interconnect to provide low latency and high messaging rate. Performance demands from the cluster interconnect increase exponentially with scale due in part to all-to-all communication patterns. This demand is even more dramatic as simulations involve greater complexity to properly simulate physical model behaviors.

In this paper we will focus on the value of In-Network computing for HDR InfiniBand Networks for OpenFOAM application for few CPU architecture.