



Contribution ID: 4

Type: **not specified**

InfiniBand In-Network Computing Technology and Roadmap

Monday, 20 September 2021 09:50 (20 minutes)

The latest revolution in HPC is the effort around the co-design approach, a collaborative effort to reach Exascale performance by taking a holistic system-level approach to fundamental performance improvements, is In-Network Computing. The CPU-centric approach has reached the limits of its scalability in several aspects, and In-Network Computing acting as “distributed co-processor” can handle and accelerates performance of various data algorithms, such as reductions and more.

The past focus for smart interconnects development was to offload the network functions from the CPU to the network. With the new efforts in the co-design approach, the new generation of smart interconnects will also offload data algorithms that will be managed within the network, allowing users to run these algorithms as the data being transferred within the system interconnect, rather than waiting for the data to reach the CPU. This technology is being referred to as In-Network Computing, which is the leading approach to achieve performance and scalability for Exascale systems. In-Network Computing transforms the data center interconnect to become a “distributed CPU”, and “distributed memory”, enables to overcome performance walls and to enable faster and more scalable data analysis.

The new generation of HDR 200G InfiniBand In-Network Computing technology includes several elements - Scalable Hierarchical Aggregation and Reduction Protocol (SHARP), a technology that was developed by Oak Ridge National Laboratory and Mellanox and received the R&D100 award, that enables to execute data reduction algorithm on the network devices instead of the host based processor. Other elements include smart Tag Matching and rendezvoused protocol, and more. These technologies are in use at some of the recent large scale supercomputers around the world, including the top TOP500 platforms.

The session will discuss the In-Network Computing technology and testing results from various systems around weather based applications.

Primary authors: SHAINER, Gilad (NVIDIA); GRAHAM, Rich (NVIDIA)

Presenters: SHAINER, Gilad (NVIDIA); GRAHAM, Rich (NVIDIA)

Session Classification: Session 1

Track Classification: 19th Workshop on high performance computing in meteorology