Contribution ID: **272** Type: **Poster**

Data Transfer Nodes and Demonstration of 100G - 400G Wide Area Throughput Using the Caltech SDN Testbed

Tuesday, 11 October 2016 16:30 (15 minutes)

We review and demonstrate the design of efficient data transfer nodes (DTNs), from the perspectives of the highest throughput over both local and wide area networks, as well as the highest performance per unit cost. A careful system-level design is required for the hardware, firmware, OS and software components. Furthermore, additional tuning of these components, and the identification and elimination of any remaining bottlenecks, is needed once the system is assembled and commissioned, in order to obtain optimal performance. For high throughput data transfers, specialized software is used to overcome the traditional limits in performance caused by the OS, file system, file structures used, etc. Concretely, we will discuss and present the latest results using Fast Data Transfer (FDT), developed by Caltech, and RFTP developed by Stonybrook together with BNL.

We will present and discuss the design choices for three generations of Caltech DTNs. Their transfer capabilities range from 40Gbps to 400Gbps. Disk throughput is still the biggest challenge in the current generation of available hardware. However, new NVME drives combined with RDMA and a new NVME network fabric are expected to improve the overall data-transfer throughput and simultaneously reduce the CPU load on the end nodes.

Secondary Keyword (Optional)

Network systems and solutions

Primary Keyword (Mandatory)

High performance computing

Tertiary Keyword (Optional)

Primary author: MUGHAL, Azher (California Institute of Technology (US))

Co-author: NEWMAN, Harvey (California Institute of Technology (US))

Presenter: MUGHAL, Azher (California Institute of Technology (US))

Session Classification: Posters A / Break

Track Classification: Track 6: Infrastructures