



Evaluation of the Pre-production Implementation of the Intel **Omni-Path** Interconnect Technology (Highlights)

Alexandr Zaytsev
alezayt@bnl.gov

BROOKHAVEN
NATIONAL LABORATORY

BNL, USA
RHIC & ATLAS Computing Facility

Summary / Highlights

- Remote evaluation of the pre-production Intel Omni-Path hardware and software was performed by the RACF in Dec 2015 using the Intel owned “Diamond” cluster provided with a non-blocking Intel Omni-Path fabric based on a single 48-port pre-production Omni-Path switch
- Performance measurements have demonstrated that it is feasible to extract close to the specification limits performance from the pre-production Intel Omni-Path hardware for the latency-bound applications utilizing the range of MPI message sizes from 1 byte to 4 MB:
 - Carrying up to **96.9 Gbps** of aggregated unidirectional MPI message traffic though a single interface (with 100 Gbps as the hardware limit)
 - Observe point-to-point unidirectional latency between two compute nodes in a single switch based fabric that is as low as **0.93 μ s**
 - Bring a single IP over Omni-Path interface up to the level of **99.2 Gbps** of aggregated unidirectional TCP/IP traffic with multi-threaded iperf tests
 - The hardware and software features of pre-production Intel Omni-Path technology make it a viable alternative to the Mellanox 4X EDR Infiniband for both latency-bound and bandwidth-bound HPC and HTC workloads