

# Evaluation of the pre-production implementation of the Intel Omni-Path interconnect technology

*Tuesday, October 11, 2016 3:15 PM (15 minutes)*

This contribution gives a report on the remote evaluation of the pre-production Intel Omni-Path (OPA) interconnect hardware and software performed by RHIC & ATLAS Computing Facility (RACF) at BNL in Dec 2015 - Feb 2016 time period using a 32 node "Diamond" cluster with a single Omni-Path Host Fabric Interface (HFI) installed on each and a single 48-port Omni-Path switch with the non-blocking fabric (capable of carrying up to 9.4 Tbps of the aggregate traffic if all ports are involved) provided by Intel. The main purpose of the tests was to assess the basic features and functionality of the control and diagnostic tools available for the pre-production version of the Intel Omni-Path low latency interconnect technology, as well as the Omni-Path interconnect performance in a realistic environment of a multi-node HPC cluster running RedHat Enterprise Linux 7 x86\_64 OS. The interconnect performance metering was performed using the low level fabric layer and MPI communication layer benchmarking tools available in the OpenFabrics Enterprise Distribution (OFED), Intel Fabric Suite and OpenMPI v1.10.0 distributions with pre-production support of the Intel OPA interconnect technology built with both GCC v4.9.2 and Intel Compiler v15.0.2 versions and provided in the existing test cluster setup. A subset of the tests were performed with benchmarking tools built with GCC and Intel Compiler, with and without explicit mapping of a test processes to the physical CPU cores on the compute nodes in order to determine whether these changes result in a statistically significant difference in performance observed. Despite the limited scale of the test cluster used, the test environment provided was sufficient to carry out a large variety of RDMA, native and Intel OpenMPI, and IP over Omni-Path performance measurements and functionality tests. In addition to presenting the results of the performance benchmarks we also discuss the prospects for the future use of the Intel Omni-Path technology as a future interconnect solutions for both the HPC and HTC scientific workloads.

## Tertiary Keyword (Optional)

## Secondary Keyword (Optional)

High performance computing

## Primary Keyword (Mandatory)

Network systems and solutions

**Primary author:** ZAYTSEV, Alexandr (Brookhaven National Laboratory (US))

**Co-authors:** HOLLOWELL, Christopher (Brookhaven National Laboratory); CARAMARCU, Costin (Brookhaven National Laboratory); WONG, Tony (Brookhaven National Laboratory); STRECKER-KELLOGG, William (Brookhaven National Lab)

**Presenter:** ZAYTSEV, Alexandr (Brookhaven National Laboratory (US))

**Session Classification:** Track 6: Infrastructures

**Track Classification:** Track 6: Infrastructures