

Performance Evaluation of Soft RoCE over 1 Gigabit Ethernet

Authors : Gurkirat Kaur, Manoj Kumar, Manju Bala

Abstract : Ethernet is the most influential and widely used technology in the world. With the growing demand of low latency and high throughput technologies like InfiniBand and RoCE, unique features viz. RDMA (Remote Direct Memory Access) have evolved. RDMA is an effective technology which is used for reducing system load and improving performance. InfiniBand is a well known technology which provides high-bandwidth and low-latency and makes optimal use of in-built features like RDMA. With the rapid evolution of InfiniBand technology and Ethernet lacking the RDMA and zero copy protocol, the Ethernet community has come out with a new enhancements that bridges the gap between InfiniBand and Ethernet. By adding the RDMA and zero copy protocol to the Ethernet a new networking technology is evolved, called RDMA over Converged Ethernet (RoCE). RoCE is a standard released by the IBTA standardization body to define RDMA protocol over Ethernet. With the emergence of lossless Ethernet, RoCE uses InfiniBand's efficient transport to provide the platform for deploying RDMA technology in mainstream data centres over 10GigE, 40GigE and beyond. RoCE provide all of the InfiniBand benefits transport benefits and well established RDMA ecosystem combined with converged Ethernet. In this paper, we evaluate the heterogeneous Linux cluster, having multi nodes with fast interconnects i.e. gigabit Ethernet and Soft RoCE. This paper presents the heterogeneous Linux cluster configuration and evaluates its performance using Intel's MPI Benchmarks. Our result shows that Soft RoCE is performing better than Ethernet in various performance metrics like bandwidth, latency and throughput.

Keywords : ethernet, InfiniBand, RoCE, RDMA, MPI, Soft RoCE

Conference Title : ICCSCN 2014 : International Conference on Communication Systems and Computer Networks

Conference Location : New York, United States

Conference Dates : June 05-06, 2014