



Contribution ID: 70

Type: **Sectional reports**

Improving Networking Performance of a Linux Cluster.

Thursday, 7 July 2016 14:45 (15 minutes)

Networking is known to be a "bottleneck" in scientific computations on HPC clusters. It could become a problem that limits the scalability of systems with a cluster architecture. And that problem is a worldwide one since clusters are used almost everywhere.

Expensive clusters usually have some custom networks. Such systems imply expensive and powerful hardware, custom protocols, proprietary operating systems. But the vast majority of up-to-date systems use conventional hardware, protocols and operating systems, for example, Ethernet network with OS Linux on cluster nodes.

This article is devoted to the problems of small and medium clusters that are often used in universities. We will focus on Ethernet clusters with OS Linux.

This topic will be discussed by the example of implementing a custom protocol. TCP/IP stack is used very often in cluster computing, even small clusters use it. While it was originally developed for Internet and could impose unnecessary overheads when it is used in a small cluster with reliable network.

We will discuss different aspects of Linux networking stack (e.g. NAPI) and modern hardware (e.g. GSO and GRO); compare performance of TCP, UDP, custom protocol implemented with raw sockets and as a kernel module; discuss possible optimizations. As a result several recommendations on improving networking performance of Linux clusters will be given.

Keywords:

Computational clusters, Linux, Networking, Protocols, Kernel, Sockets, NAPI, GSO, GRO.

Primary author: Mr GAIDUCHOK, Vladimir (Saint Petersburg Electrotechnical University "LETI", Russia)

Co-authors: Prof. BOGDANOV, Alexander (St.Petersburg State University); Mr GANKEVICH, Ivan (Saint Petersburg State University); Mr AHMED, Nabil (Saint Petersburg Electrotechnical University "LETI", Russia); Mr IVANOV, Pavel (Saint Petersburg Electrotechnical University "LETI", Russia)

Presenter: Mr GAIDUCHOK, Vladimir (Saint Petersburg Electrotechnical University "LETI", Russia)

Session Classification: 3. Middleware and services for production-quality infrastructures

Track Classification: 3. Middleware and services for production-quality infrastructures