



Contribution ID: 548

Type: **Poster**

NUMA memory hierarchies experience with multithreaded HEP software at CERN openlab

Tuesday, 22 May 2012 13:30 (4h 45m)

Newer generations of processors come with no increase in their clock frequency, and the same is true for memory chips. In order to achieve more performance, the core count is getting higher, and to feed all the cores on a chip with instructions and data, the number of memory channels must follow the same trend. Non Uniform Memory Access (NUMA) architecture allowed the CPU manufacturers to reduce nicely the impact of memory subsystem bottlenecks, but, in turn, this solution introduces a cost at the application level. This paper describes our practical experience with the typical CPU servers currently available to the HEP community, based on work with NUMA systems at CERN openlab. We provide the latest measurements of the different NUMA implementations from AMD and Intel, as well as NUMA consequences on some parallelized HEP codes.

Primary authors: Dr LAZZARO, Alfio (CERN openlab); NOWAK, Andrzej (CERN); LEDUC, Julien; Mr JARP, Sverre (CERN)

Presenter: LEDUC, Julien

Session Classification: Poster Session

Track Classification: Computer Facilities, Production Grids and Networking (track 4)