

CPU capacity metric for WLCG

26 May 2007 – Manfred Alef, Helge Meinhard, Les Robertson

*The proposals in this document were approved by the
WLCG Management Board – 5 June 2007*

The LHC experiment Computing TDRs, the WLCG TDR and the WLCG MoU all use thousands of SPECint2000 (KSI2K) units as the metric for CPU capacity. At the time these documents were prepared (2005) the SI2K benchmark results published for different processors on the spec.org site matched fairly closely the relative performance of high energy physics codes on computer systems using these processors. Over the past two years, however, there has been an increasing divergence between published SPECint numbers and HEP program performance, noted especially in the case of the recent multi-core processors.

On the other hand, some WLCG sites (e.g. FZK, CERN) have used an alternative methodology for running the SPEC benchmark. Operating system, compiler and compilation options are chosen such that they resemble closely the way CPU servers are used in daily life, as many copies of the benchmark as there are processor cores in the shared memory system are executed simultaneously and the system performance is defined to be the sum of the individual benchmark results. This has a significantly better correlation with the observed performance of HEP programs.

WLCG sites have reacted to this problem in different ways. Some sites run their own benchmarks and install equipment and report usage according to their own estimates of system performance. Other sites simply use the manufacturer's SPECint rating. The discrepancy between these measures can be in excess of 50%

In order to enable experiments to plan their usage of current and planned CPU resources it is necessary to re-establish a base for processor performance. A working group on benchmarking has been established recently by HEPiX and we expect that it will in due course make proposals for a new standard for HEP. In the meantime the following interim solution is proposed.

- We adopt the methodology of GridKA (see the [HEPiX Benchmarking page](#) at CASPUR) to measure system performance in KSI2K units. The compiler flags used will be specified from time to time by the Architects Forum, with the initial set that defined as “*low optimization*” on the HEPiX page.
- We then increase this measure by **50%**, which approximates the difference between this measure and the published spec.org values in the first half of 2005.

Sites may perform these measurements themselves, or may use measurements made by other HEP sites, or by reference to the results published on the CASPUR benchmark Web site.

The new metric should be implemented by sites for reporting installed capacity and usage **by September 2007**.