

Number of Threads per Copy in HEP Workloads

Manfred Aef / HEP-SCORE Deployment Task Force 2021-03-03

Reasonable Default for “threads_per_copy”

- In typical WLCG environments, multicore jobs will get 8 cores (slots)
(2 sockets*SMT=enabled*even_no._of_cores_per_chip ⇒ multiple of 8)
- There are also (few) machines with 12, 20, or 28 cores
(SMT off, only 1 socket, ...) – multiples of 4, but not multiples of 8
- (Other special cases like 6 or 10 cores?)
- HEPiX Benchmarking Working Group had suggested default of 4 threads per copy for the multithreaded/multiprocessing workloads to occupy the whole machine in workload runs
- What’s the impact (e.g. memory consumption, swapping) when running multicore workloads with 4 threads per copy instead of 8?
- Testing all valid combinations with 4...10 threads per copy on hosts with 40, 48, and 256 cores/logical processors

Results of WL Runs

CPU	2x AMD EPYC 7702 64-Core Pr. (256 threads)		2x Intel Xeon E5-2630v4 (40 threads)				4x AMD Opteron 6174 (48 cores) (2.0 GB RAM per core)		
	Copies	64	32	10	8	5	4	12	8
Threads p. copy	4	8	4	5	8	10	4	6	8
ATLAS Sim	.502	.487	.082	.081	.079	.078	.060	.059	.057
ATLAS Digi Reco	25.85	19.14	2.91	2.89	2.95	2.72	2.83	2.82	2.68
CMS Gen Sim	5.85	5.89	.943	.952	.953	.953	.781	.821	.778
CMS Digi	26.90	28.22	4.37	4.44	4.40	4.37	3.35	3.48	3.38
CMS Reco	18.35	18.34	2.71	2.71	2.71	2.71	2.20	2.23	2.18
ATLAS Gen LHCb Gen Sim Belle	single-core workloads								

Summary

- No significant differences in the scores when running 4 or 8 threads per copy (differences of up to some 5%)
 - *Outlier: ATLAS Digi Reco when running on host with many cores (here: 2x AMD EPYC 7702 with SMT enabled (2*64*2=256 logical processors))*
 - No swapping detected when running the workloads on machines with not more than 2 GB per core
-
- **Proposal by HEPiX Benchmarking Working Group:
Stay with default setting of 4 threads per copy**