




ChEESE

Partner: BSC - CNS



**Barcelona
Supercomputing
Center**
Centro Nacional de Supercomputación


Marta García
Victor López



Project funded by EuroHPC under the grant agreement No 101093038.



Task 2.4 Status Update (BSC)





Task 2.4 Status Update (BSC)

- DLB pre-release (3.6.0-beta) supporting GPU metrics for CUDA & AMD
- Release 3.6.0 planned for Nov'25 (this week)
 - C, Fortran, Python
 - MPI, OpenMP, OpenACC, CUDA, HIP
- Links of interest:
 - <https://pm.bsc.es/dlb-downloads>
 - <https://github.com/bsc-pm/dlb>
 - <https://dlb-docs.readthedocs.io/en/latest/>

Task 2.4 Status Update (BSC)

- Metrics for CUDA devices working as expected
 - `DLB_ARGS="--talp --plugin=cupti"`
- Metrics for AMD devices working on LUMI, needs for testing on other platforms
 - requires `--plugin=rocprofilerv2` for ROCM < 6.2
 - requires `--plugin=rocprofiler-sdk` for ROCM >= 6.2

FALL3D

Table 2. TALP Output for FALL3D from 1 to 8 nodes

Metrics		Nodes			
		1	2	4	8
Host	Parallel efficiency	0.26	0.16	0.10	0.07
	└ MPI Parallel eff.	0.44	0.27	0.16	0.11
	└ └ Comm. eff.	0.85	0.83	0.82	0.89
	└ └ Load Balance	0.52	0.32	0.20	0.12
	└ Device Offload eff.	0.59	0.61	0.63	0.64
Device	Parallel efficiency	0.14	0.08	0.04	0.03
	└ Load Balance	0.98	0.97	0.96	0.96
	└ Communication eff.	0.78	0.77	0.76	0.75
	└ Orchestration eff.	0.19	0.11	0.06	0.04

FALL3D




Fig. 13. FALL3D Traces showing useful durations. Top: 1 node, 4 GPUs. Bottom: 8 nodes, 32 GPUs.

XSHELLS

Table 3. TALP Output for XSHELLS from 1 to 8 nodes

Metrics		Nodes			
		1	2	4	8
Host	Parallel efficiency	0.36	0.29	0.26	0.15
	└ MPI Parallel eff.	0.90	0.64	0.51	0.25
	└ └ Comm. eff.	0.91	0.66	0.52	0.27
	└ └ Load Balance	0.98	0.97	0.98	0.93
	└ Device Offload eff.	0.40	0.45	0.51	0.60
Device	Parallel efficiency	0.52	0.35	0.24	0.10
	└ Load Balance	1.00	1.00	1.00	1.00
	└ Communication eff.	0.98	0.97	0.96	0.94
	└ Orchestration eff.	0.54	0.36	0.25	0.10

XSHELLS



Fig. 14. XSHELLS traces showing CUDA calls. Top: 1 node, 4 GPUs. Bottom: 4 nodes, 16 GPUs