



Motivation

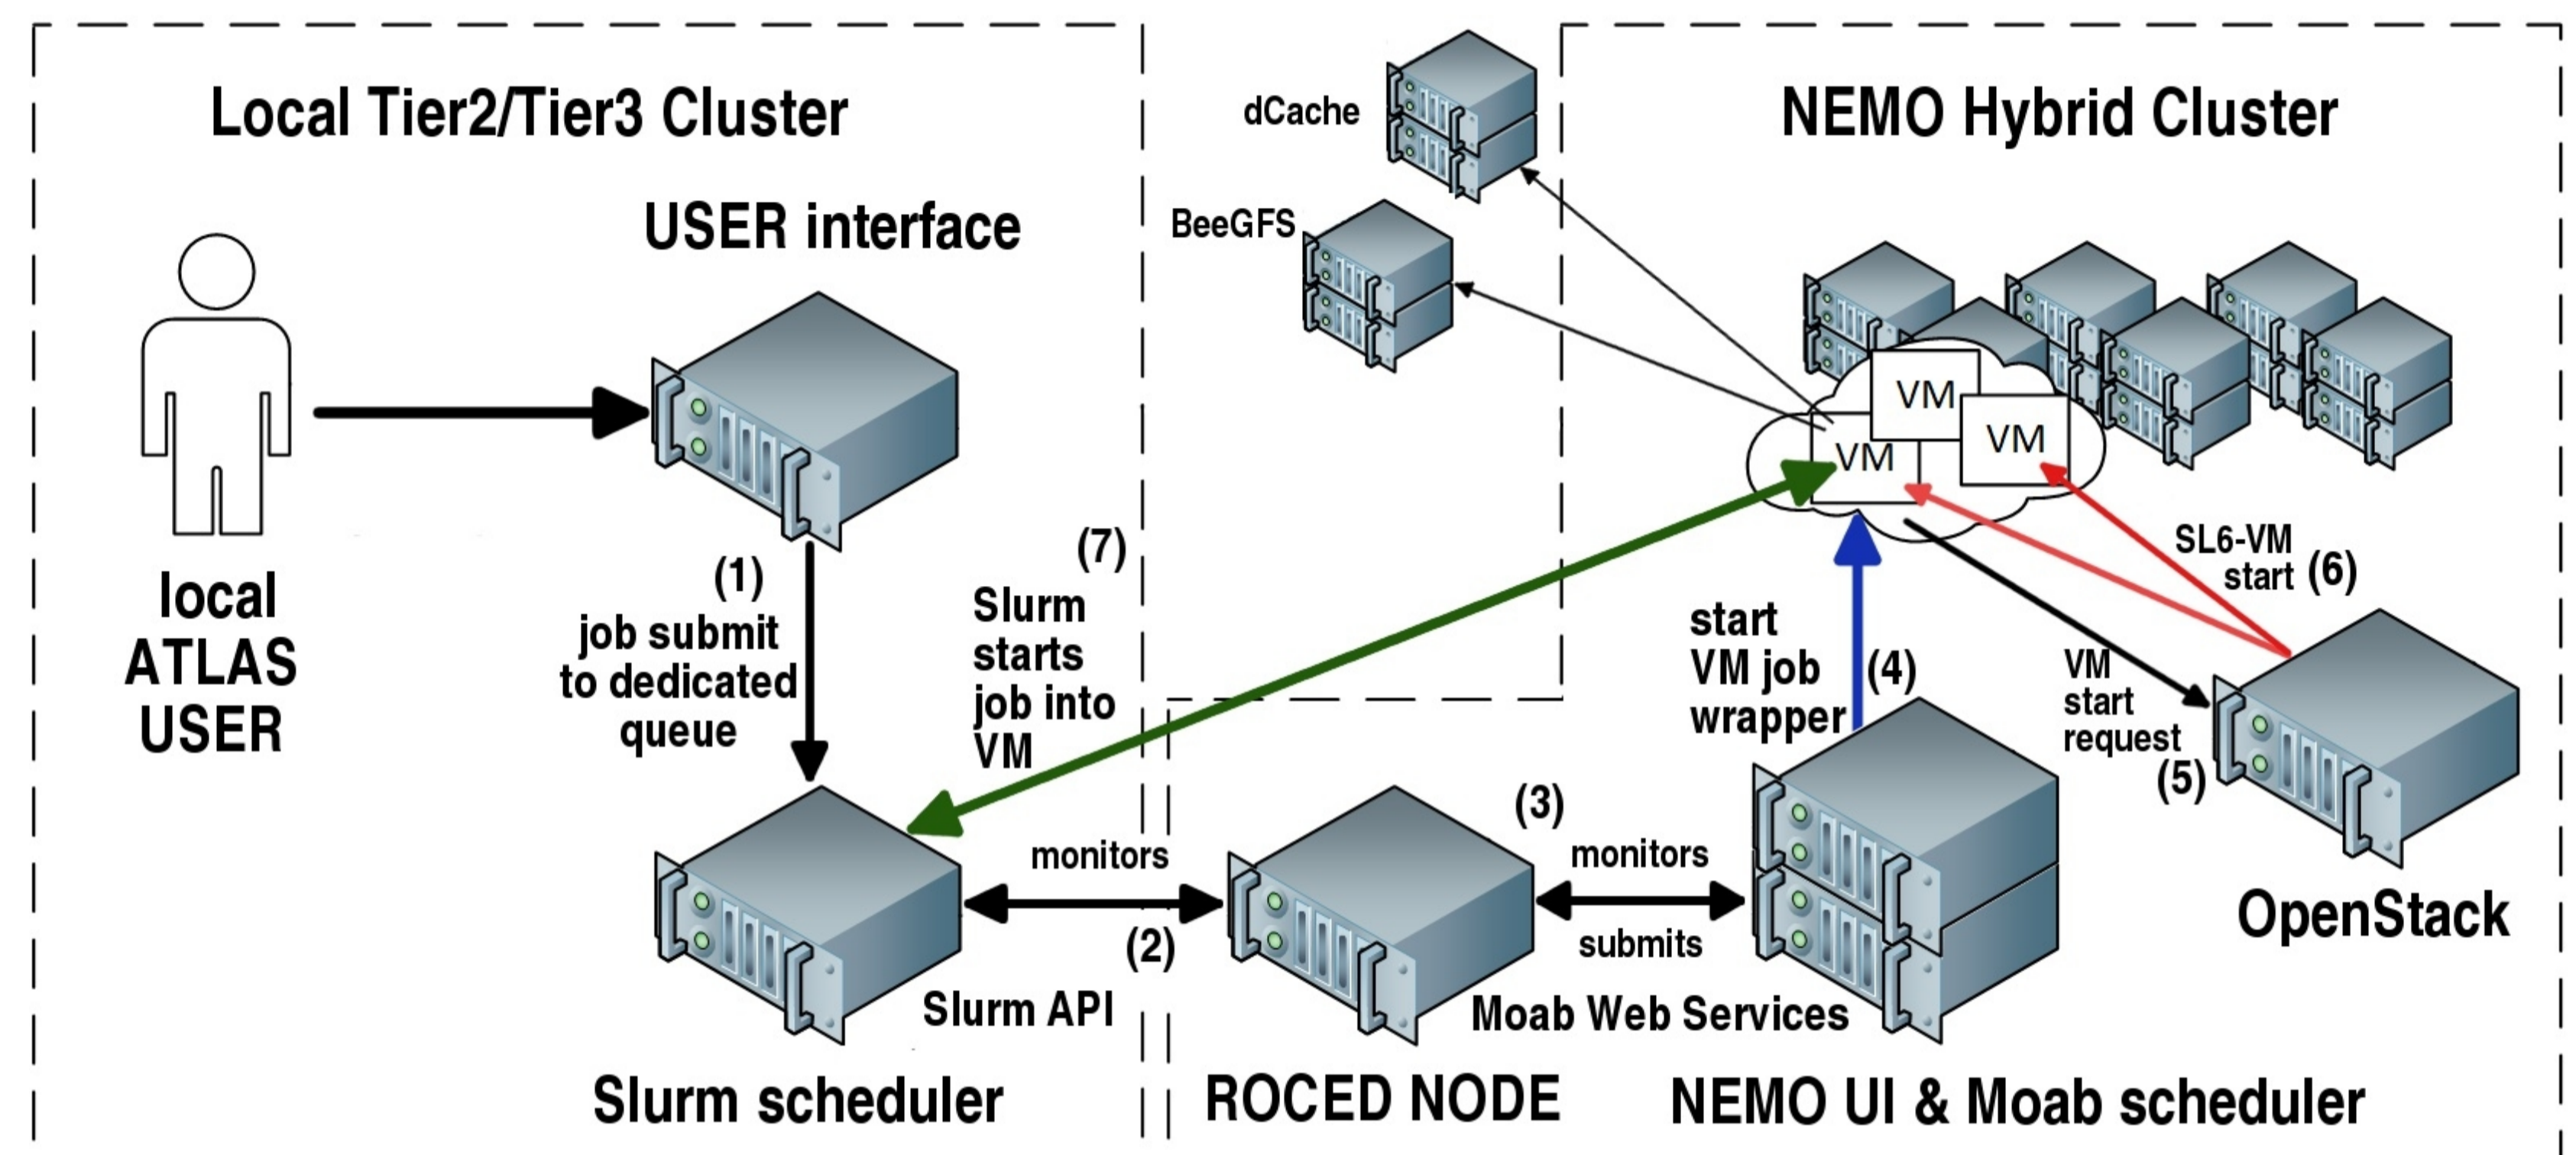
- Use University HPC cluster NEMO, CentOS7, 324 000 HEP-SPEC06
- Develop a model to integrate opportunistic resources without admin privileges

Tasks

- Provide full analysis and production environment for local users (Tier3) and for ATLAS / WLCG (Tier2) jobs
- Easy and automatic provisioning of VMs
- Monitor demand of local Tier3 cluster
- Start VMs on-demand via NEMO scheduler and OpenStack
- Evaluate performance of VMs in setup

Batchsystem Integration

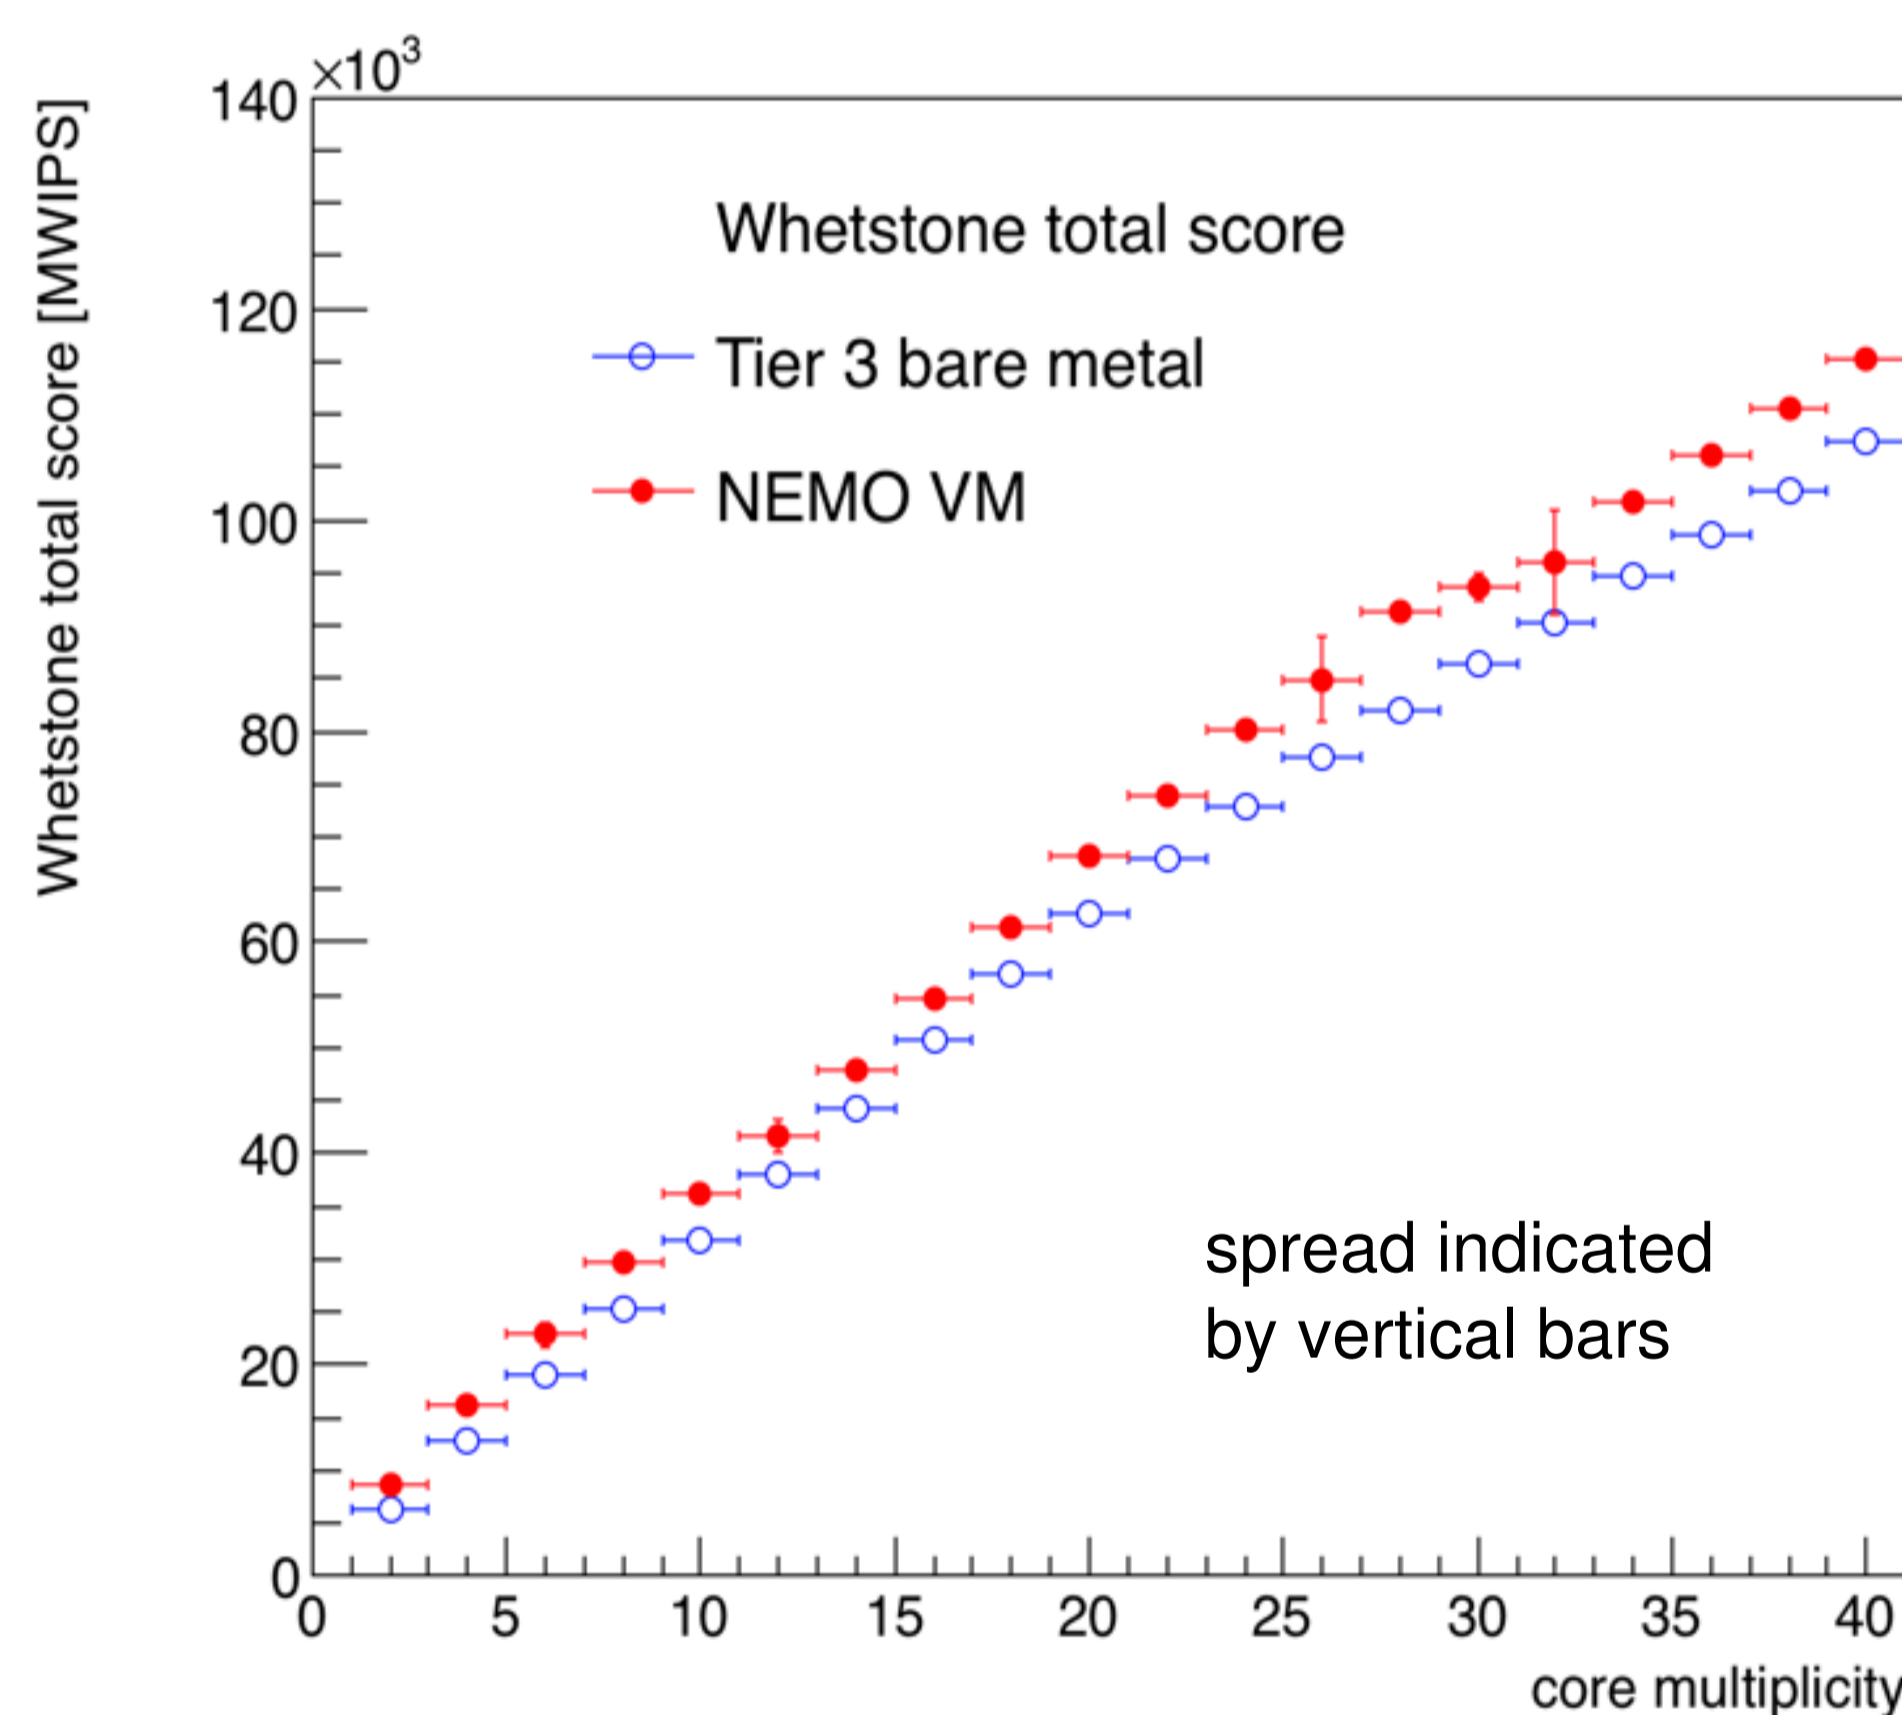
- Fully contextualized provisioning of VM (Virtual Machines) by Packer and Puppet
- ROCED coordinates Slurm and Moab on-demand
- Job wrapper requests VM start from OpenStack
- Test readiness of VM before setting online



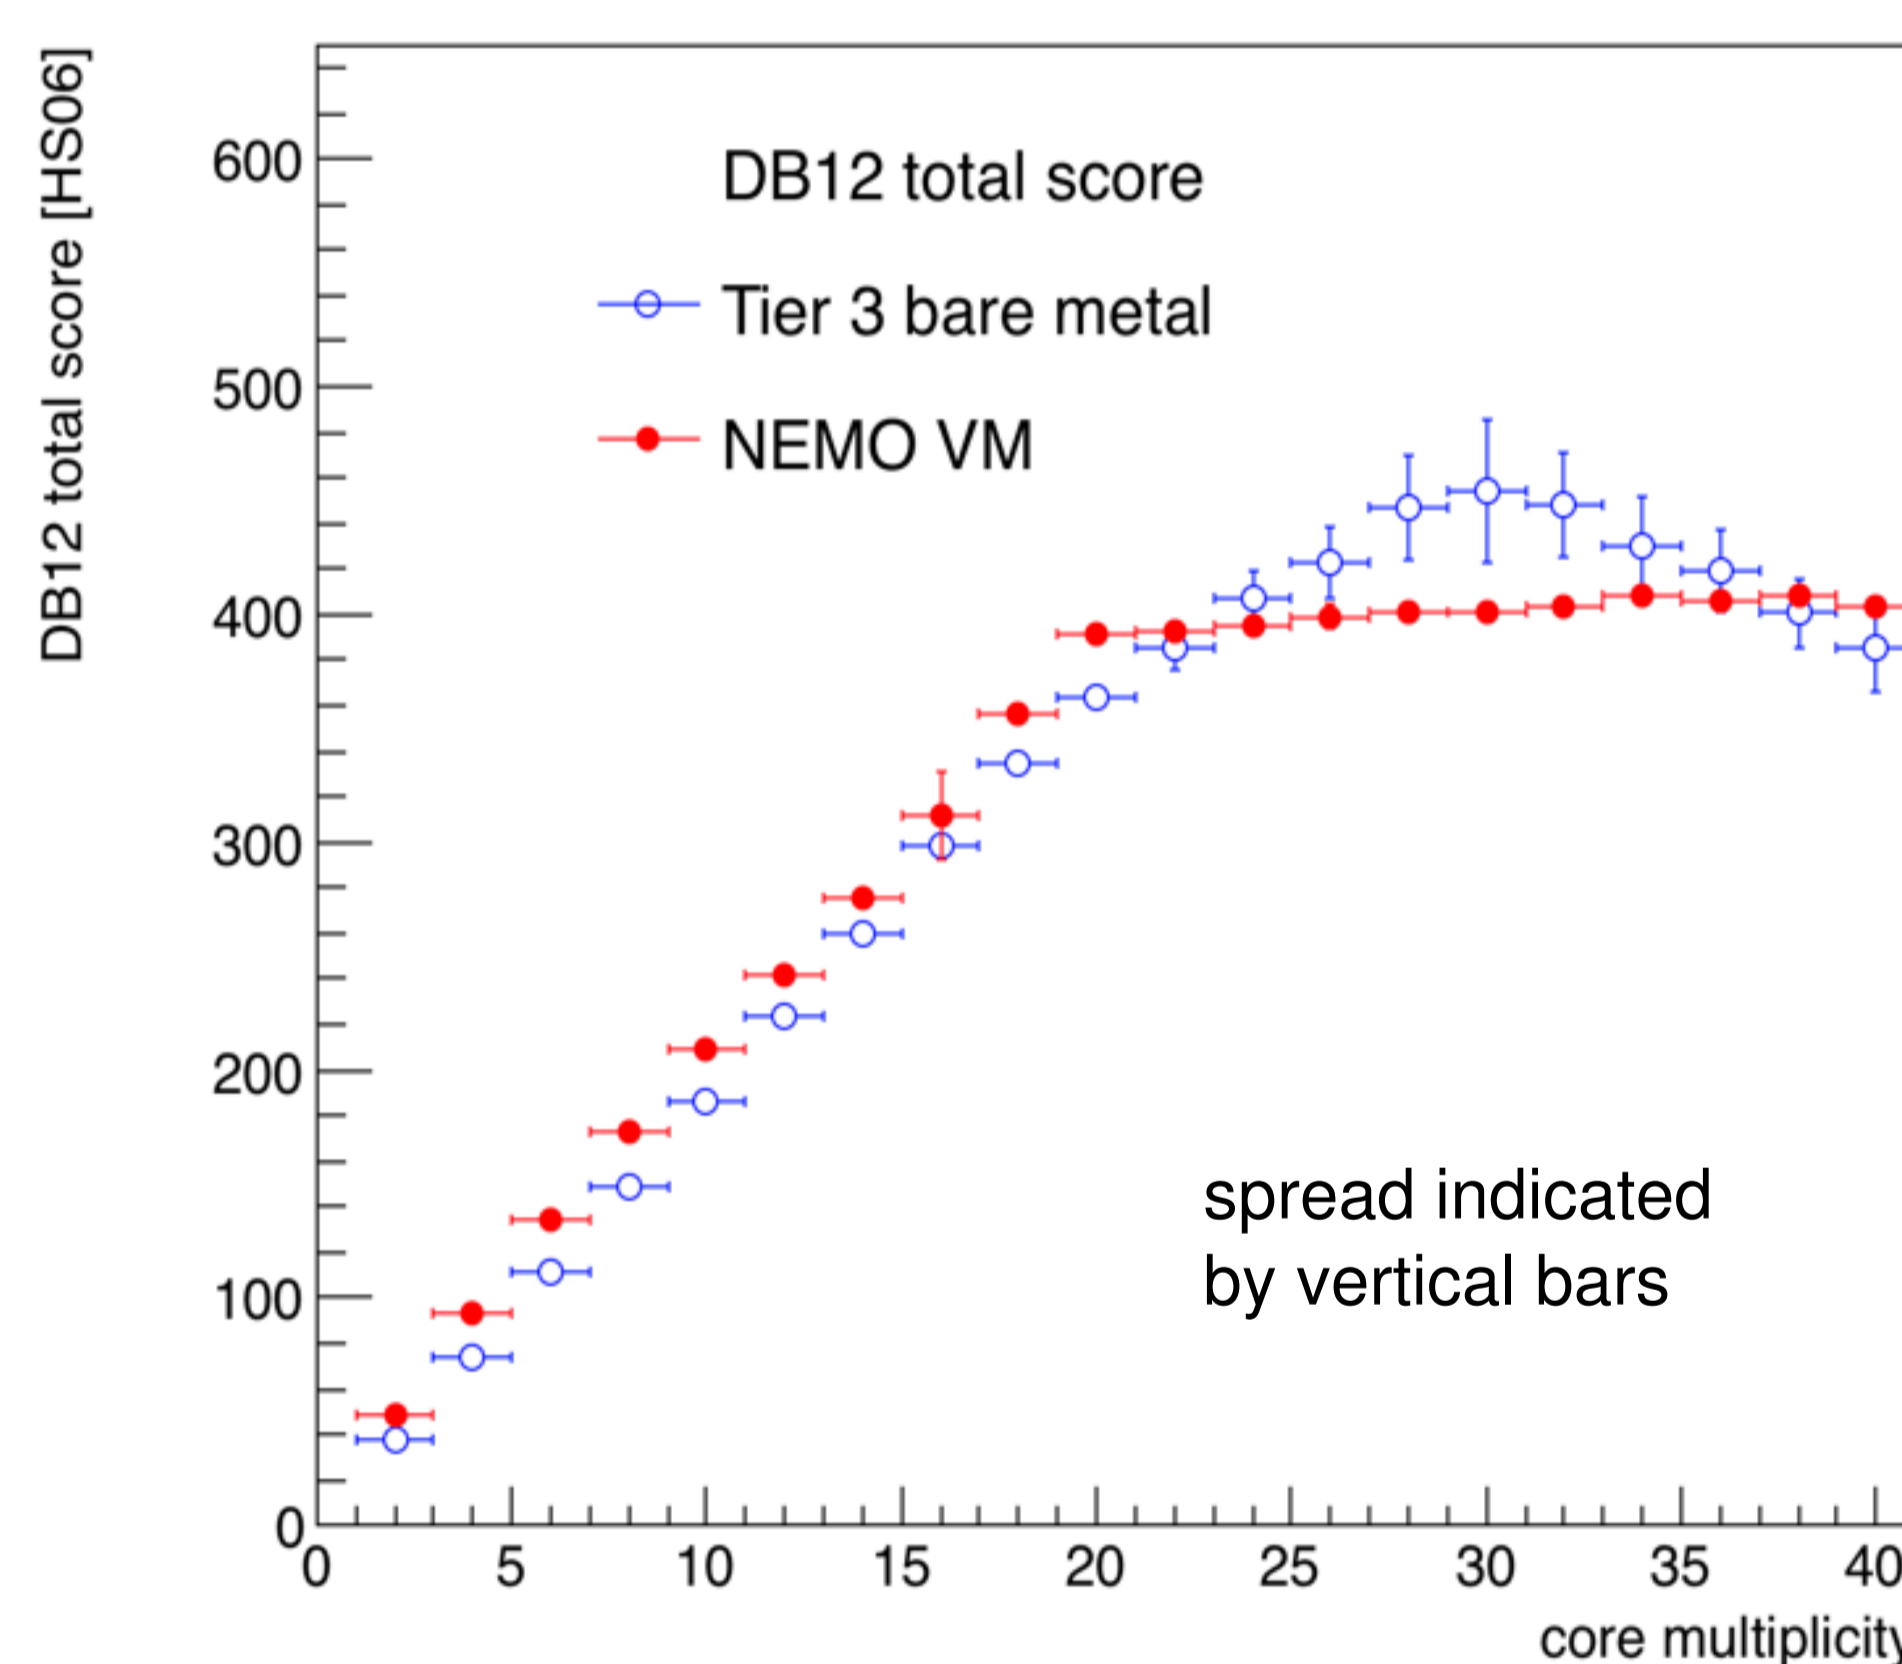
Performance Evaluation

- Compare performance with standard HEP-SPEC06 benchmark and new benchmarks
- Using DB12 (Dirac benchmark) and Whetstone (MWIPS) from CERN benchmark suite to get quick reference of performance
- Hardware: All tests on 2x INTEL CPU E5-2630v4 2.20GHz 40 cores HT on INTEL S2600KPR board, 128GB RAM DDR4
- CentOS7 host (20-core limit), Tier3 SL6 diskless install, bare metal, multicore vs. SL6 VM image (40-core) on NEMO

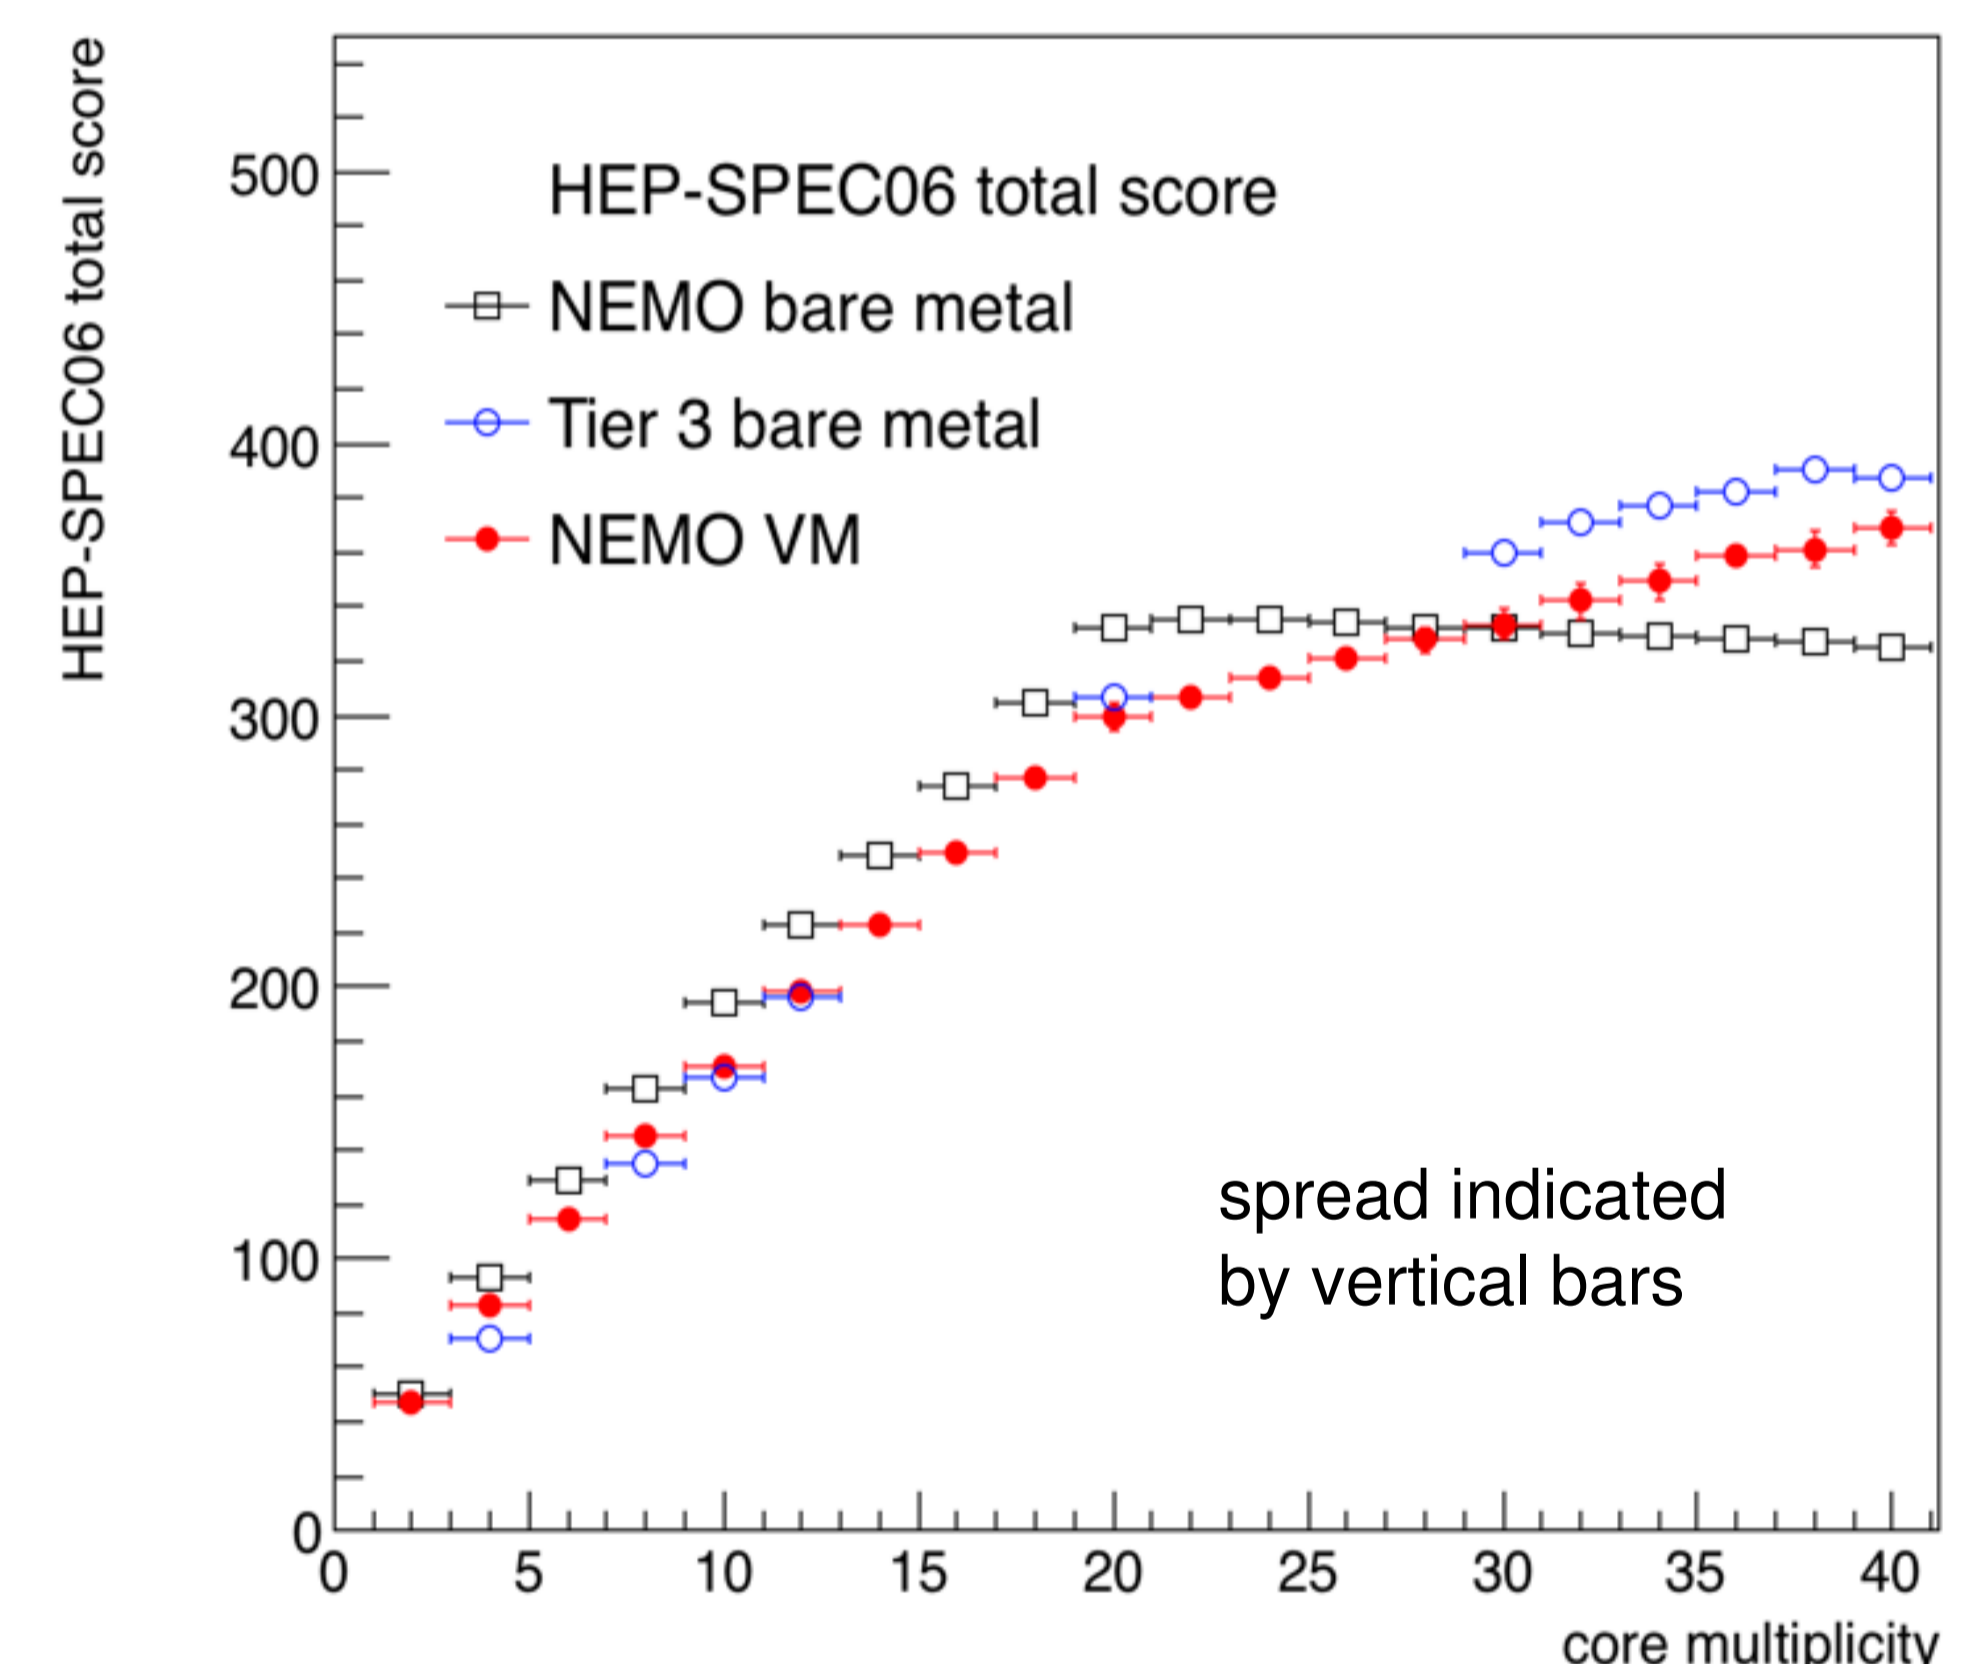
1) Whetstone benchmarks
elapsed time: 2 – 3 min



2) DB12 Dirac benchmarks
elapsed time: 4 – 6 min



3) HepSPEC 2006 benchmarks
elapsed time: 2 – 6 hours



Whetstone benchmarks show a linear gain of performance also on HT cores

DB12 Dirac benchmarks: linear gain for cores <20 ; some fluctuations for the bare metal HT cores >20

No big loss of performance in VMs compared to bare metal on different operating systems.

Results

- ROCED integrates local Tier2/Tier3 Slurm and NEMO Moab by supervising both schedulers
- SL6 VMs on NEMO opportunistic resource show comparable performance to native SL6 / CentOS7
- In addition to HEP-SPEC06 DB12 and Whetstone are useful quick benchmarks in our setup

Outlook

- Integrate CERN benchmarks as part of Hammercloud tests to evaluate performance of nodes on-site
- Also use other application benchmarks e.g. from ROOT to evaluate performance
- Develop monitoring and analytics systems to further improve ROCED