Rethinking Compute Platforms: When Cloud Can Replace HPC, and When It Cannot

M. Foglieni*, E. Birang, P. Dobrev, B. Emrem, P. Ganta, F. Jamitzky

Leibniz Supercomputing Centre (LRZ), 85748 Garching, Germany

*matteo.foglieni@lrz.de

Traditional HPC (tHPC) challenges:

- Long deployment time (weeks or months)
- Limited flexibility for short-term projects
- Complex software stack management
- Hard reproducibility across different centers

IaC-based virtual HPC (vHPC) advantages:

- Rapid deployment: production-ready clusters in hours
- Reproducibility: version-controlled infrastructure configurations
- Resource elasticity: scale up for burst workloads, scale down when idle
- HPC-prototyping: testing new technologies for GCS (LRZ, JSC, HLRS)
- Lower barriers: simplified access for educational purposes
- Simple software management: EESSI (European Environment for Scientific Software Installations)

Methodology

vHPC Deployment: Magic Castle [1]

- Terraform-based automated cluster deployment on OpenStack with 4 Nodes
- Pre-configured SLURM workload manager
- Integrated software stack (EESSI)
- Customization: Spack, Conda Environments, Open OnDemand installation

Hardware setup

Specs	vHPC	tHPC
Name	GCS DevCloud (InHPC-DE project)	SuperMUC-NG Phase 1
Number of nodes	4	6338
CPU	Intel IceLake Xeon Silver 4316	Intel SkyLake Xeon Platinum 8174
Physical Cores per node	40	48
Total number of cores	160	304128
Memory chan/speed	8 / 2667 MHz	6 / 2666 MHz
RAM per node	980 GB	96 GB
Interconnect	Ethernet 100 Gbit/s	OmniPath 100Gbit/s

Key Findings

Network Latency is the Primary Bottleneck

Virtualization overhead stems primarily from inter-node communications. Furthermore, the network hardware connection in the GCS cloud are Ethernet, which is not suited for communication-intensive workloads.

Application-Specific Scaling Behavior

Quantum ESPRESSO and LAMMPS shows near-ideal scaling on the vHPC, while GROMACS and CP2K exhibit stronger sensitivity to the network bottleneck.

Conclusions

IaC-based virtual HPC enable much faster, reproducible and easily customizable deployments of HPC-like systems, with great performance for low-communicating workflows compared to traditional HPC.

Network latency is the primary optimization target, both on software (virtualized layer) and on hardware (Ethernet vs Infiniband), but even with such limitation vHPC is production-ready for education, prototyping, and burst computing scenarios.

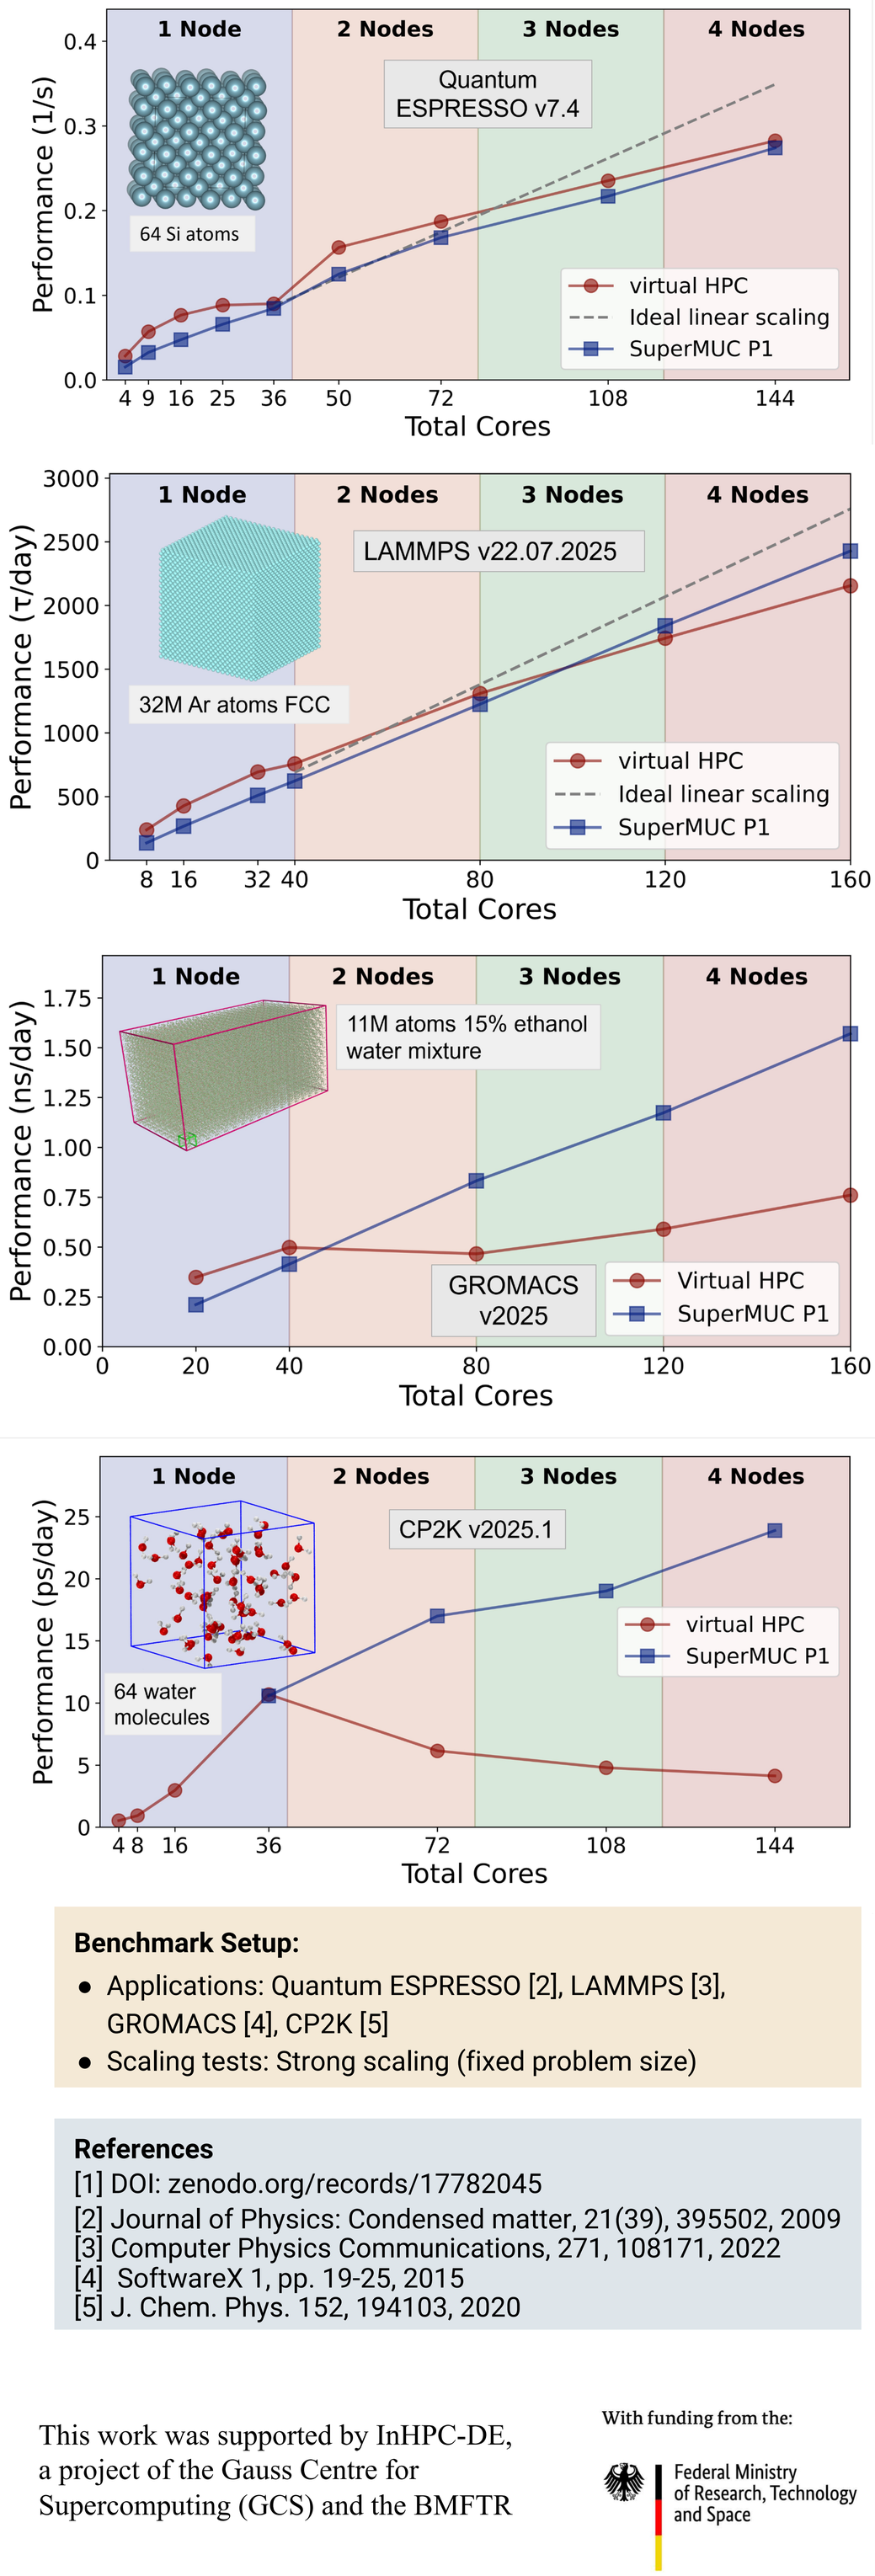

Benchmark Setup:

- Applications: Quantum ESPRESSO [2], LAMMPS [3], GROMACS [4], CP2K [5]
- Scaling tests: Strong scaling (fixed problem size)

References

[1] DOI: zenodo.org/records/17782045
[2] Journal of Physics: Condensed matter, 21(39), 395502, 2009
[3] Computer Physics Communications, 271, 108171, 2022
[4] SoftwareX 1, pp. 19-25, 2015
[5] J. Chem. Phys. 152, 194103, 2020

Leibniz Supercomputing Centre
of the Bavarian Academy of Sciences and Humanities

Gauss Centre for Supercomputing

This work was supported by InHPC-DE, a project of the Gauss Centre for Supercomputing (GCS) and the BMFTR

With funding from the:
Federal Ministry of Research, Technology and Space