

Cloud In Cresco (CLIC): towards HPC/HPDA-as-a-service

Marco Faltelli

ENEA

HPC current challenges

- Users' accessibility
 - Specialized knowledge is often required (packages, HW...)
 - Lack of specific libraries
- Clusters are fragmentated, each one has its own platform
 - ENEA example: CRESCO4, CRESCO5, CRESCO6, CRESCO7, CRESCO8...
- Maintenance
 - Causes partial/total downtime periods
- Failures
 - Especially in long-lasting jobs

In-cloud HPC: goals

- Unified access to resources
 - Unique login portal for all clusters
 - The orchestrator chooses where to allocate VMs
 - New clusters can be added modularly
- Specialized environments:
 - E.g., custom VMs for AI with packages and UI
 - High-level platform for non-skilled users (SaaS-like)
- Increased resilience:
 - For users: snapshots and live migration in case of failure
 - For administrators: easier maintenance through AZs and migration

CLIC preliminary performance tests

Openstack infrastructure, Ubuntu 22.04, virtual GPU and virtual Infiniband cards

NVIDIA HPC (A100-40G)		
	vGPU (98 SM)	bare-metal (108 SM)
HPL 21.4	9.73 TFLOPS	9.85 TFLOPS
HPCG 24.09	228 GFLOPS	230 GFLOPS

Infiniband (Mellanox 100 GbE)		
	VM	bare-metal
OFED perf (read)	96 Gbps	96 Gbps
OFED perf (write)	96 Gbps	96 Gbps
MPI host-to-host	88 Gbps	88 Gbps

Upcoming cluster with 24 H100 GPUs in Brindisi

Context and future goals



Funded by
the European Union
NextGenerationEU



Ministero delle Imprese
e del Made in Italy

- This work is funded by the 8ra (IPCEI-CIS) initiative
- The proposed infrastructure is a PoC for HPC in a cloud-edge continuum
- CLIC will be distributed among 4 ENEA DCs
- Sites upgrade ongoing, bid process is next

- Special thanks to: A. Peloso, F. Iannone, M. Fois, M. Celino, and G. Ponti
- **For further infos:** marco.faltelli@enea.it