



SubMIT: Hardware Resources and Performances

Mariarosaria D'Alfonso

Annual Review on Basic computing Services May. 22, 2025



SubMIT





very flexible system, users utilize the sw/hw according to their requirements

I will focus on hardware resources and performance

- \rightarrow disk resources, compute, network
- \rightarrow status, capacity usage
- \rightarrow benchmarking of the system



Computing resources allocation (cpu and gpu)



20% of users also use GPU



Data Storage allocation (/home /work /scratch)



/home	5GB User's home with backed-up storage notebooks and local code developments
/work	50GB for software installations
/scratch	NVMe disk with fast access (for short term storage) ultrafast
Ceph	1TB per user and groups space to store larger datasets large-scale high-performance

/work is 63 TB (12% used) /home is 70 TB (3% used) /scratch is 44 TB (full) ... *more needed ? ...*





Data Storage allocation: ceph



ด ceph

1TB per user and groups space to store larger datasets

 \rightarrow Users can request additional storage space or a group storage of reasonable size

• If they request more space, we usually ask their group to contribute some storage disks

 \rightarrow We have 826 individual ceph storage, and 126 of our users make use of this space





Ceph usage per individual per division





Ceph usage per group



Astro 4.0 ABCP NUPAX NUPAT 3.5 3.0 Number of groups 2.5 -2.0 -1.5 1.0 0.5 0.0 + 10^{-1} 100 10¹ 10² 10³ 104 105 106 Usage (GB)

Ceph Group Storage Usage in GB



2024-2025 Enhancement



hw resources integrated:

1. computing:

- a. added 1 new machine with 384 cores (PPC submit82)
- b. 2 machines 2x36 cores (lain Stewart /CTP submit95-96)
- c. 2 machines with 4 Nvidia V100 to subMIT (CMS, previously at Bates)
- 2. disk: 17 x 22 TB of spinning disk on cephFS (+30% since last year)
- 3. network:
 - a. Links from 10Gbps to 100Gbps
 - b. LACP channel bonding: Simultaneously use of multiple ethernet connections for single servers (2x100Gbps) (2x100Gbps) and (6x1Gbps)

infrastructures improvements:

we have capacity for new machines within our retirement policies and the space given new powerstrips added





Benchmarking: I/O



Run a benchmark analysis on CMS-open data:

Used a powerful machine:

384 threads, 1.5 TB of RAM, 100Gbit/s link, run RDF with intrinsic multi-threading

Performance comparison across storage system



NVMe /scratch connected with 100 Gbits/s links \rightarrow fastest read

xcache (ESNet site in Cambridge, NVMe) \rightarrow fast read

remote read EOS FNAL/CERN are fast

HDD cephFS /submit & T2 affected by the minimum read size 16MB as configured when using buffered IO, much more than we need



Benchmarking: distributed computing



Run a benchmark analysis on CMS-open data:

Benchmark the distributed computing (RDF + DaskGateway + Slurm) on subMIT:

read data from the /scratch NVMe disk connected with 100Gbit/s links

Performance comparison across computing configurations







SubMIT is a powerful and user-friendly working system already enabling high-impact scientific research and education.

Its seamless integration of <u>high-performance computing</u>, <u>high-speed networking</u>, and <u>high-throughput and large mass storage</u> forms the backbone of SubMIT.

The current 2025 resources demonstrate strong capabilities, with clear provision for future scalability and growth.