



HPC Data Challenge



David Southwick, Maria Girone, et al

IT-GOV-INN

Openlab HPC efforts

CERN Openlab partners with external industry & organizations

HPC sites host major concentrations of GPU/accelerator HW

- All Quantum access will be via HPC (at least in Europe)

HPC adoption efforts from CoE RAISE (Openlab EC project), FZJ Jülich, RTU, BSC

- Snowball effect: Expanded to interTwin, Fenix, *DOMA* (thanks Mario!)
- HPC site operators are aware of upcoming big-data science needs, they want to prepare


HPC Data Challenge

Substantial activity ramp-up from experiments over past years in HPC adoption

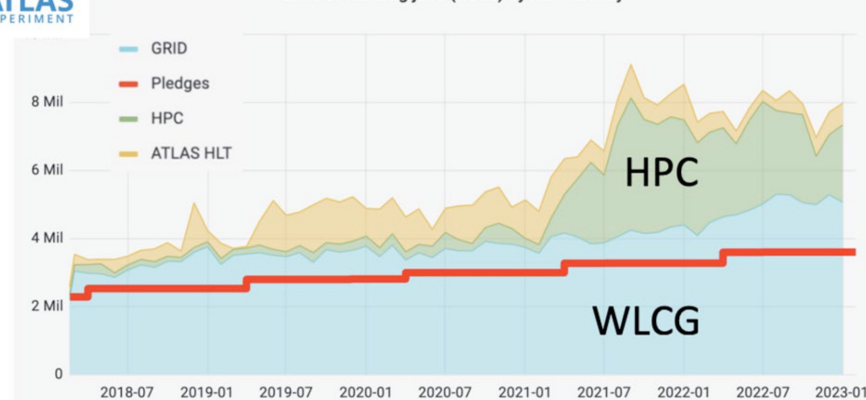
- experiments (soon) fulfilling pledges with HPC resources

CMS Public

Number of Running CPU Cores on HPCs - Monthly Average



Slots of Running jobs (HS06) by ADC activity



HPC (mini) Data Challenge

Structure HPC data challenge similar to DOMA; challenge of increasing complexity

2029: ~10PB data through a HPC site in a day *

- Initial 10% goal (1PB), scaling steps following DataChallenge format
- demonstrate management of data, transfer tooling (**new** to HPC)
- Maintain compute efficiency with high throughput

Efforts in progress

Initial testing already underway with two HPC centers:

- FZJ (Jülich), DE
 - 200Gbps (25GB/s) peering via Géant
 - Debugging with iperf, rucio tests spring '24
- SDSC (UCSD), USA (Diego)
 - 600Gbps (75GB/s) peering via ESnet
 - Testing xrood <->FNAL + edge caching
 - long-haul tests spring '24
- Leverage Géant DTNs with “out of band” HPC sites in EU (summer '24)
- Adoption of WLCG transfer tooling stack (FENIX, ESCAPE)

Challenges

Visibility of co-related/co-located efforts (Discovering some of them from this event 😊)

High “on-boarding” cost for integrating new HPC sites (time/effort intensive)

No “batch <-> SLURM” scheduler consensus

- Several individual efforts (dask, HTcondor extensions...)
- Data locality metadata critical for “JIT”/opportunistic compute

Common framework for accessing HPC resources (SPECTRUM project '24)

-> On the path towards common HPC interface for big-data science

-> Looking for interested collaborators! Please reach out!

Thank You!



Data Lakes

Separation of WLCG sites responsibilities to new “Data Lake” model for LHC data storage has introduced new standards and modernized capabilities. Leveraging better data access patterns to datasets with latency-hiding advancements of XrootD/Xcache greatly reduces data transfer requirements:

- RUCIO – a high level data management layer, coordinates file transfers over several protocols (HTTP/WebDAV, XrootD, S3, etc.)
- FENIX – Collaboration of HPC sites and ESCAPE to standardize data transfers

