


WLCG grid and HPC: site requirements

Laurent Duflot (IJCLAB LCG-FR)




Overview



ARC Computing Element

ARC 7




Includes data movement via disks local to ARC-CE



Readonly filesystem of unique immutable files, centrally published and made accessible through cache layers.

Locally accessible of Worker Node via fuse mount and local cache

Usually used for distribution of software, config files, low volume data





Production applications (1)

◆ Generator(s)

- ◆ Simulates a collision in terms of particles produced
- ◆ No data as input
- ◆ No LAN/WAN connection needed
- ◆ Compute intensive
- ◆ Some generator can run multi-core

◆ Simulation

- ◆ Simulates energy deposition in the detector
- ◆ Take generator output as input (file ~small)
- ◆ Can be made to run without LAN connection using special containers
- ◆ Compute intensive



Production applications (2)

◆ Digitisation

- ◆ Simulates the electronics response of the detector, i.e. produce data as it would come out of the experiment
- ◆ Using fat container might not need access to DB (LAN+caches to WAN) **TBC**
- ◆ ATLAS: adding spurious collisions (“pileup”) using additional input files (i/o intensive!) but may switch to the CMS solution
- ◆ CMS: spurious collisions centrally “pre-mixed” but need WAN access



Production applications (3)

- ◆ The rest of the chain runs on simulated and real data inputs but **for real data will need access to remote DB (at least LAN + caches to WAN)**
- ◆ **Reconstruction**
 - ◆ Extract information in terms of physics objects suitable to physics interpretation
 - ◆ Compute intensive
- ◆ **Derivation**
 - ◆ Make a smaller dataset suitable for analysis, applying best calibrations etc...
 - ◆ More i/o intensive than reconstruction



- ◆ Not discussed here
 - ◆ More i/o intensive
 - ◆ Very diverse in terms of software used (analysis frameworks + user code, pure user code, ML code,...) and so i/o intensity
- ◆ There is a need to access GPUs e.g. for ML training but is the grid context the easiest for this use case?
- ◆



WLCG and HPC in Europe (**ATLAS**)

- ◆ The many ways of ATLAS to make use EuroHPC machines :
 - ◆ The HPC machines of the Nordic consortium **VEGA** (Slovenia) were configured to be “grid-compatible” : ATLAS uses ARC-CE with data transfer ; CVMFS is available ; “any” kind of jobs can be ran.
 - ◆ **KAROLINA** et **BARBORA** (Czech Republic) only have ssh access to storage and batch submission. Files are exchanged with a grid storage via **sshfs**. Whole node allocations only. CVMFS is not installed but since user namespaces are configured **cvmfsexec** is used. For multi-core production jobs.
 - ◆ **LUMI** (Finland) only have ssh access to storage and batch submission. Files are exchanged with a grid storage via **sshfs**. CVMFS is not installed and user namespace not configured. ATLAS runs “fat containers” that have code + data.
 - ◆ **MARENOSTRUM5** (Spain) and **SuperMUC** (Germany) have no outside connectivity, files are transferred to/from grid storage via scp. CVMFS is not installed and user namespace not configured. ATLAS runs “fat containers” for simulation since they don’t need access to remote DB



Questions, discussion

