


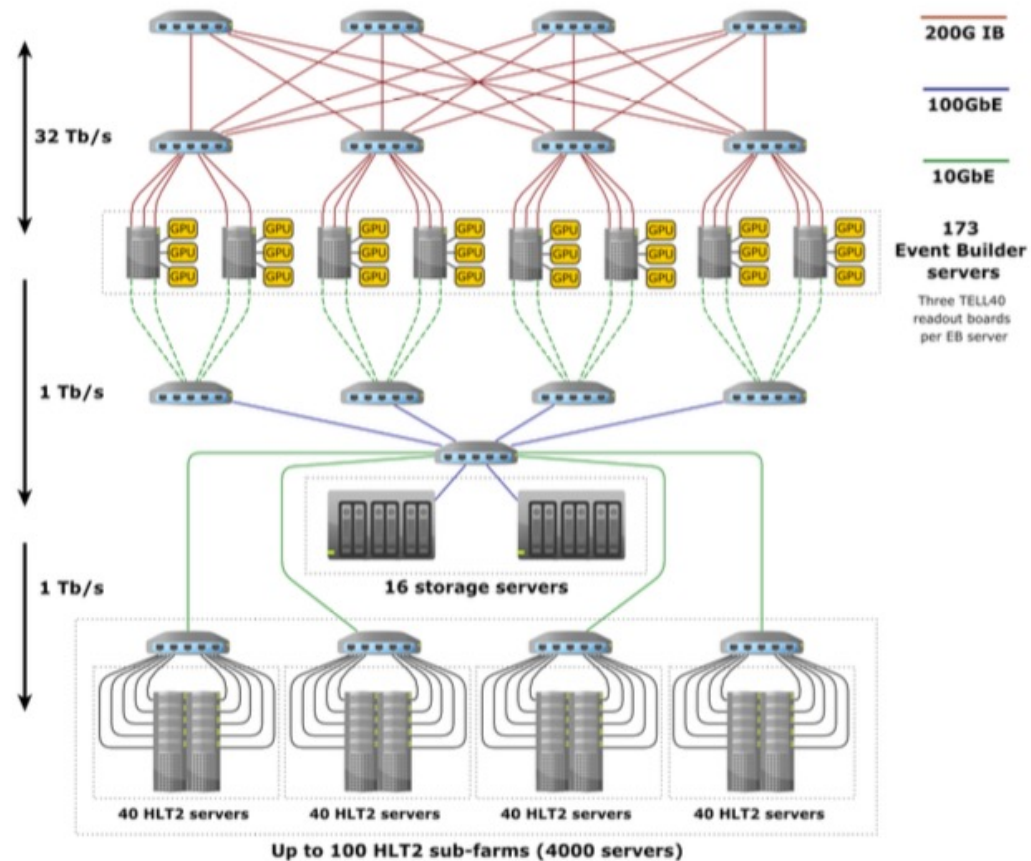
LHCb R&D activities

Concezio Bozzi

IT R&D Advisory Group Meeting

September 22nd, 2021

HLT1 on GPUs (Allen)


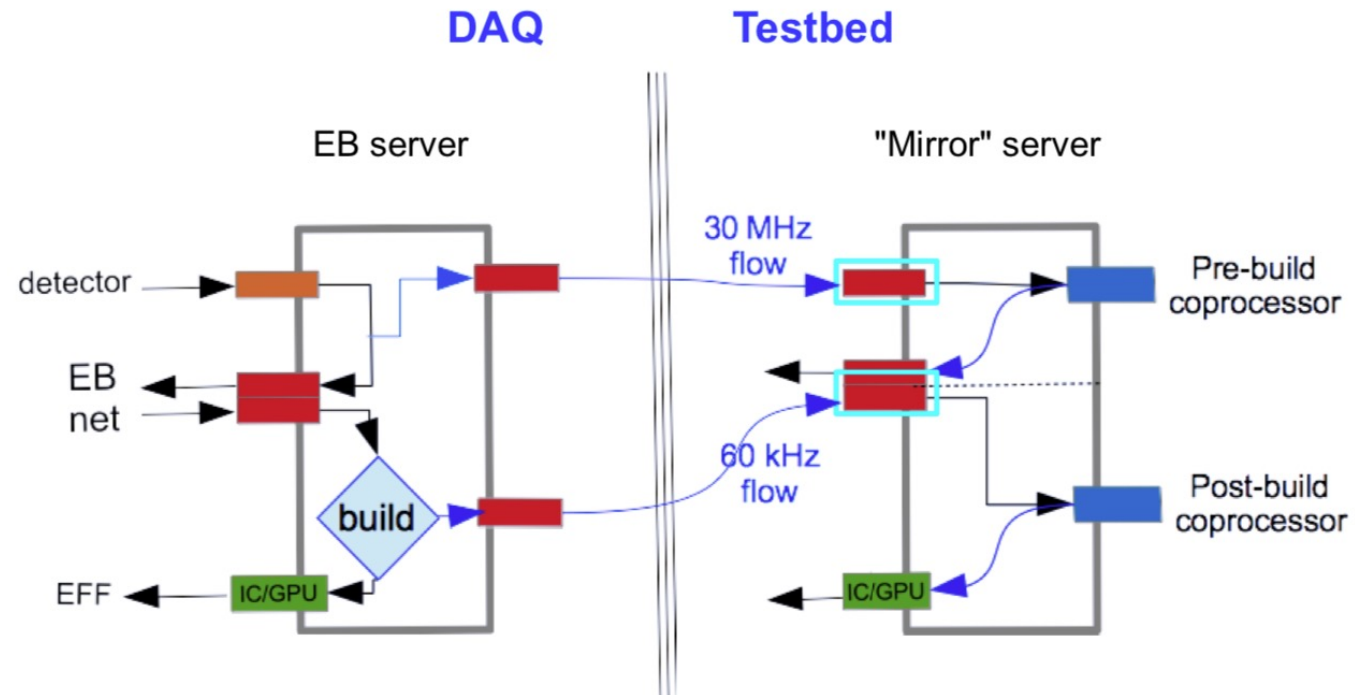


Production Readiness Review in August

- 163 Event Builder nodes
- Each will host one GPU card
- To reach 30MHz (with some margin), need to purchase high-end cards
- Chosen GPU: Nvidia A5000


Coprocessor testbed

- Operate coprocessors (of any kind) in close vicinity of the Event Builder
- Either on pre-build, or post-build data
 - Physically inside EB servers or nearly so
- Setup reproducing the hardware, software, and data flow environment of the LHCb DAQ
- With the ability to see parasitically the real data flowing through the DAQ during data taking



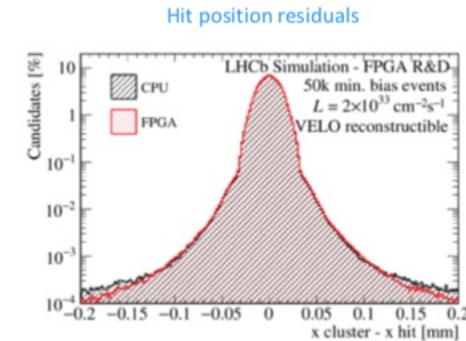
Coprocessor testbed

Second generation IPU systems for AI infrastructure at scale



- IPU could be used for highly parallel fitting
 - currently being tested to perform Kalman-filter tracking
- Retina: cellular automata designed for track reconstruction
 - Use it e.g. to reconstruct downstream tracks upfront in Run 4 on FPGAs.
 - In Run 3: Velo clustering on the Tell40 at no additional cost → increasing HLT1 throughput

Retina Cluster finder for the VELO in Run3



- In the pipeline:
 - FPGA-based NN for VELO tracking
 - Dedicated calorimetry clustering
 - Serenity boards

Non-x86 CPUs

- Validating LHCb software stack to ARM
 - LHCb environment built/distributed for ARM with conda
 - LCG stack is [available for ARM](#)
 - Using a ThunderX node provided by CERN IT/Openlab
- Investigating PPC
 - Pilot project with Italian LHC groups to access CINECA / Marconi 100 resources in Bologna (GPU+Power9)
- **Q: Can non-x86 R&D servers be provided as a service?**
 - i.e. like to OpenLab ARM server, but which are stable in production and can be used, e.g. to run nightlies and/or PR jobs.
 - important to have access to whole node, possibly without virtualisation

Machine Learning on GPUs

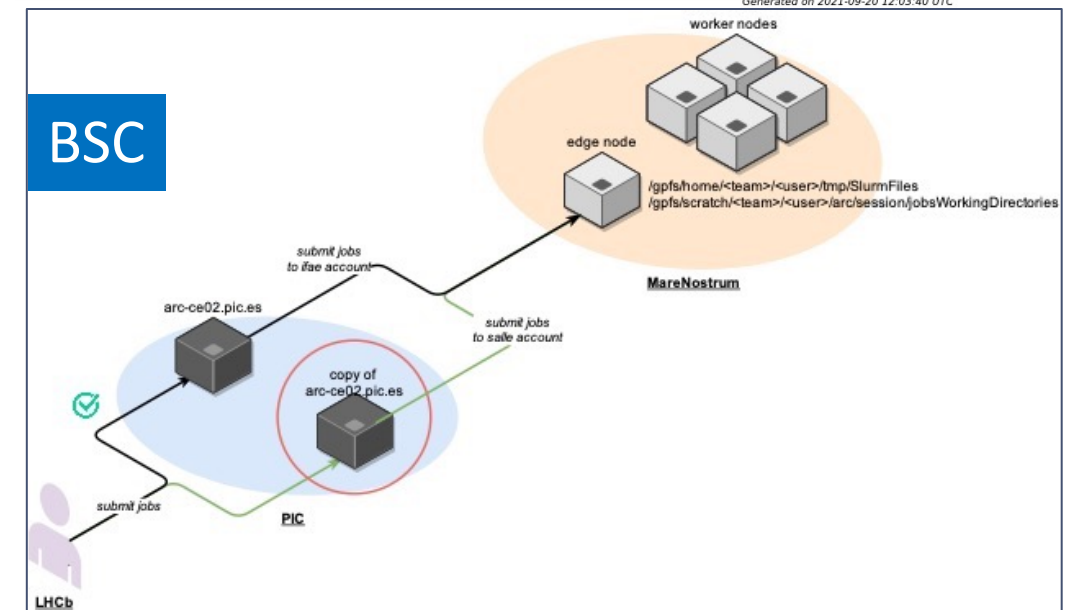
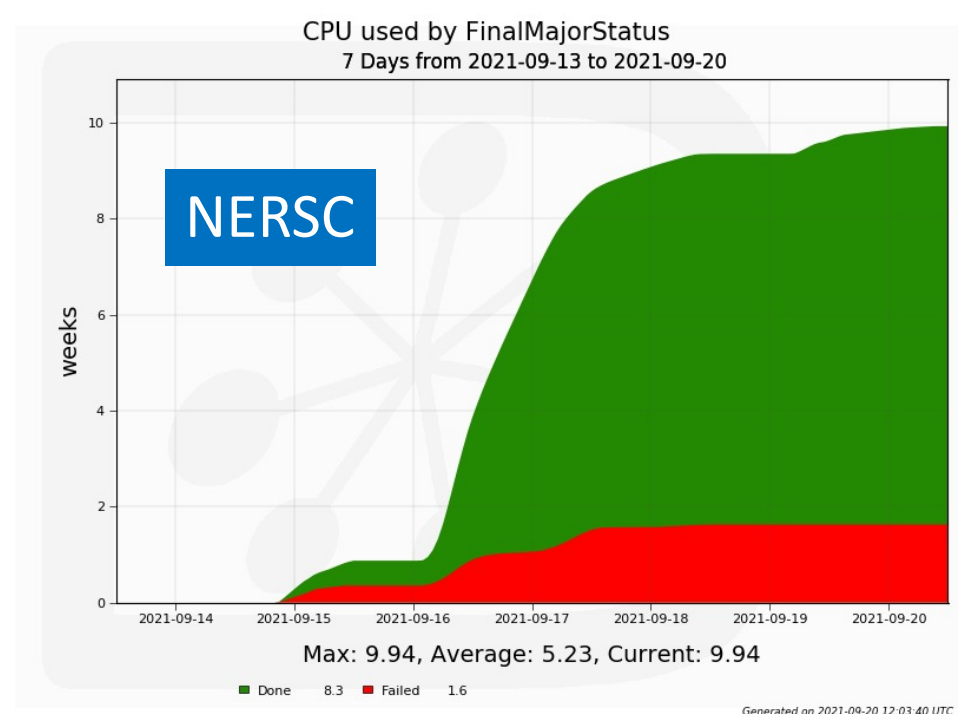
- Various activities can benefit from CERN IT/OpenLab resources and additional procurements (e.g. Cloud Broker Pilot)
- **Training of GANs** to replace the simulation of parts of the detector with deep generative models
 - Use large datasets of calibration samples to provide a model of detector response
 - Several tens of training cycles required for hyperparameter optimization
- Deployment of **fast simulation based on ML algorithms**
 - Transfer the detector response from calibration samples to physics channels
 - Assess the flexibility of the resources in terms of libraries and configurations
- **Fitting large datasets** in many dimensions (e.g. zfit / GooFit)
- Training and validation of **ML-based reconstruction algorithms**
 - Use Convolutional Neural Networks (CNN) to identify primary vertices
 - Obvious synergy with HLT1 on GPUs
- Ballpark estimate of required resources: a few 10,000 GPU.hours

Setting up HPCs

- **NERSC (US DoE):**
 - A few million core hours secured by the LHCb Los Alamos group
 - Jobs running since a few days!
- **Barcelona Supercomputing Center (BSC)**
 - Waiting for the installation and configuration of an **ARC CE**
 - Support from LaSalle and ATLAS BCN groups
- **CINECA/Marconi100**
 - GPU + Power9: difficult to use in normal production workflows
 - Pilot project with Italian LHC groups
 - User jobs are run locally
 - Started to configure DIRAC for grid-like access

Take-home message:

- Each HPC is different and requires dedicated effort to commission.
- Only the CPU component can be used in production
- Porting to other platforms requires effort



Outlook

- Using in-house resources for GPU HLT1 and additional, online-related workflows
- ARM server provided by IT/OpenLab essential for building, testing, developing, benchmarking LHCb applications
 - LCG ARM release highly beneficial (thanks!)
 - Looking forward to PPC
 - Could these R&D servers be provided as a service?
- Simulation, reconstruction and analysis activities based on ML with GPUs can benefit from IT/OpenLab facilities
- HPC access needs some streamlining
 - CERN-wide concerted effort to have standardised access to HPCs?
 - without each experiment having to reinvent the wheel on each of them
- We welcome (specialised) effort
 - for helping experiments deploy on HPCs and/or for porting software stacks to non x86
 - to optimise the software for these heterogeneous platforms