



# Towards the use of HPC for HEP workflows

Maria Girone, CERN IT and CERN openlab CTO

**EXDCI-2 Virtual Workshop on AI and HPC convergence, 26.11.2020**

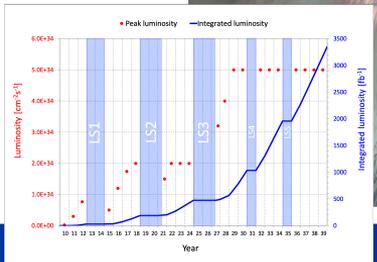
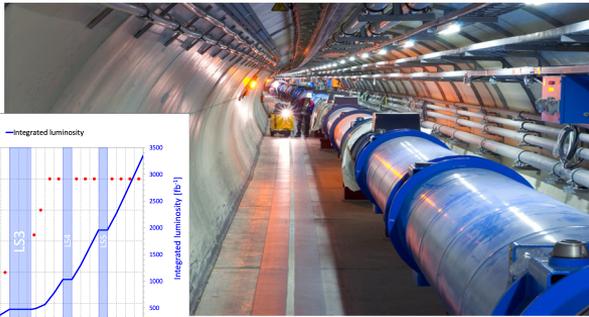
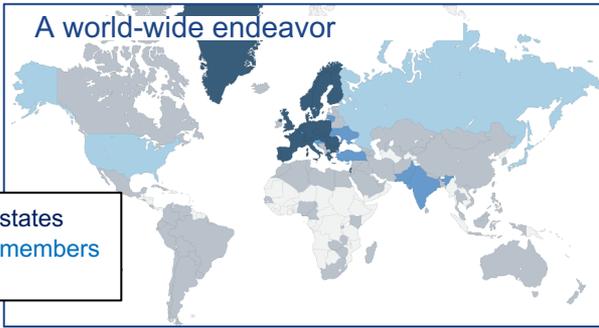
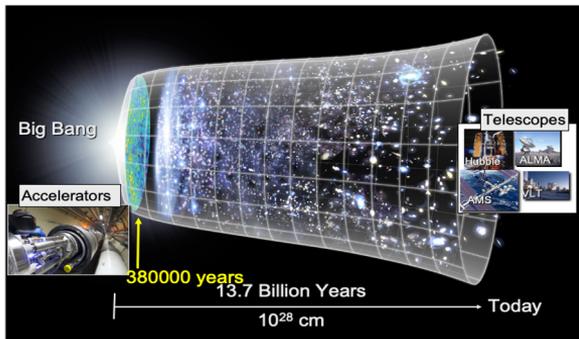
# CERN's primary mission is SCIENCE

Physicists and engineers at CERN use the world's largest and most complex scientific instruments to study the basic constituents of matter – fundamental particles

Founded in 1954, the CERN laboratory sits astride the Franco-Swiss border near Geneva

The instruments used at CERN are purpose-built particle accelerators and detectors

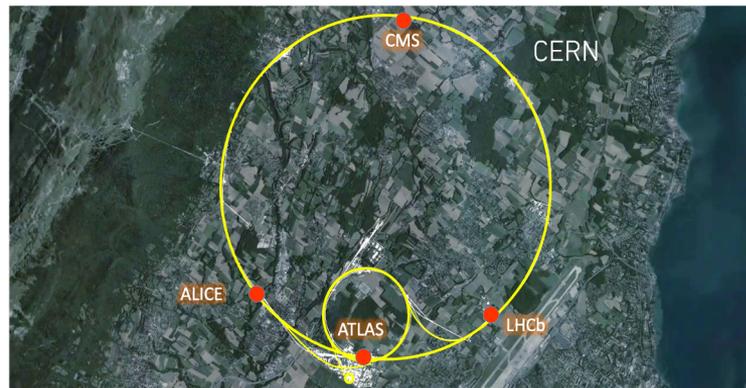
- The **Large Hadron Collider (LHC)** is the world's largest and most powerful particle accelerator and consists of a 27-kilometre ring of superconducting magnets.



# The CERN LHC Program

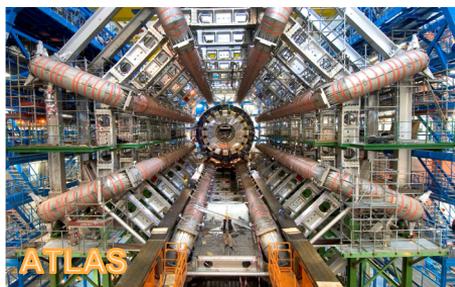
## The LHC accelerator:

- 1232 high-tech superconducting magnets
- magnet operation temperature: 1.9 K (-271 °C)
- Particles travelling at 99.9999991% the speed of light
- Beams collide: 40 million/s
- collision “temperature”:  $10^{16}$  K



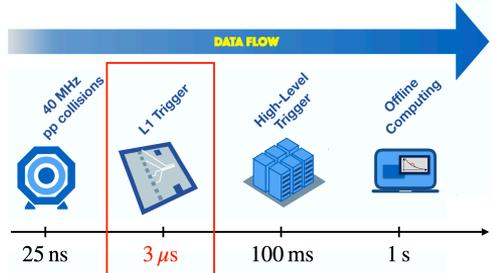
## Four large detectors at the LHC:

- size of ATLAS: ~ half Notre Dame cathedral
- weight of CMS experiment: ~13000 tons (twice the Eiffel Tower)
- number of detector sensitive elements: ~100 million
- cables needed to bring signals from detector to control room: ~3000 km
- data in 1 year: ~100 PB



# The Scale of the LHC Computing

Data processing and analysis drives physics discovery

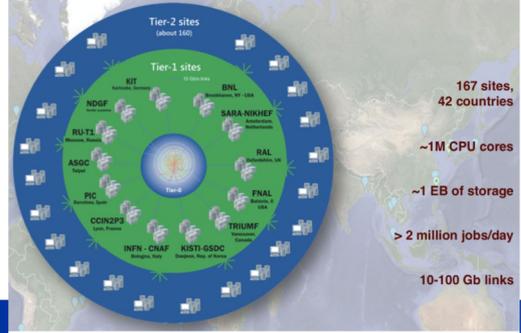
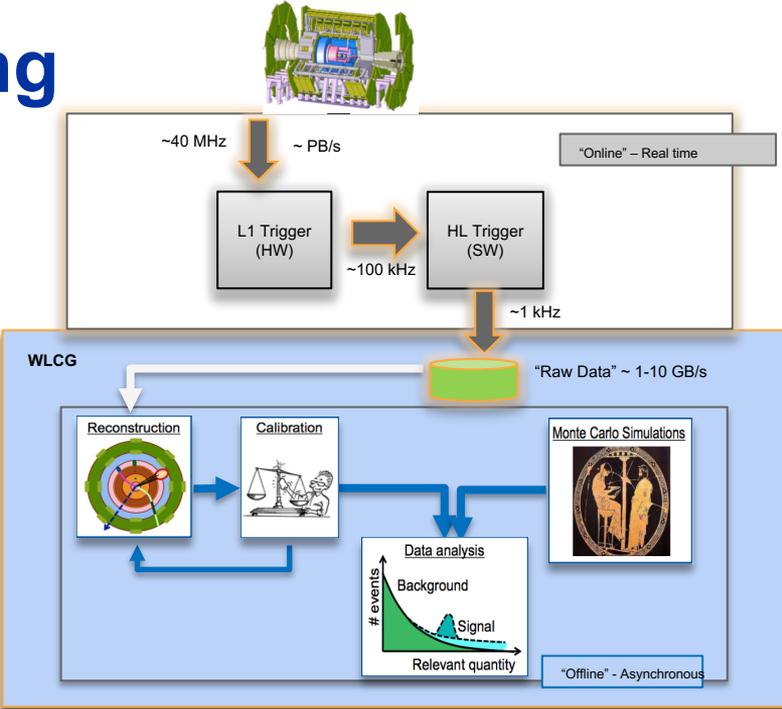


## Software

- 50M lines of C++
- Contributions by thousands of scientists
- Optimized for platforms that are evolving slowly
- Development work porting to accelerators

## Hardware

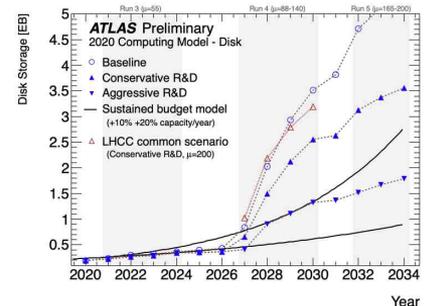
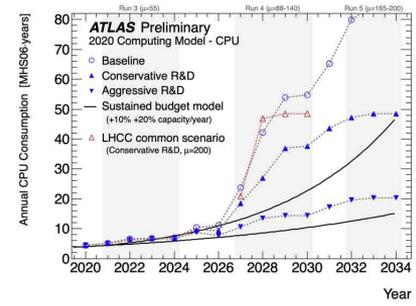
- Primary computing resources is the WLCG is a globally distributed storage and processing infrastructure
  - 167 sites over 42 countries
  - **~1M CPU cores** and **~1 exabyte** of storage (disk and tape)



# Challenges at HL-LHC



- High-Luminosity LHC will deliver about x10 increase in luminosity over LHC design
  - Pileup of 200 means a great increase in event complexity
  - More collisions and more complex data will result in a compute challenge at exascale level



# Moving to HPC

HPC falls at the intersection of several important R&D areas

- Engagement with the HPC community can be a catalyst for progress

The HPC centers and the data intensive science are both growing to exascale

- An exascale program for an exascale problem

HPC represents an opportunity to

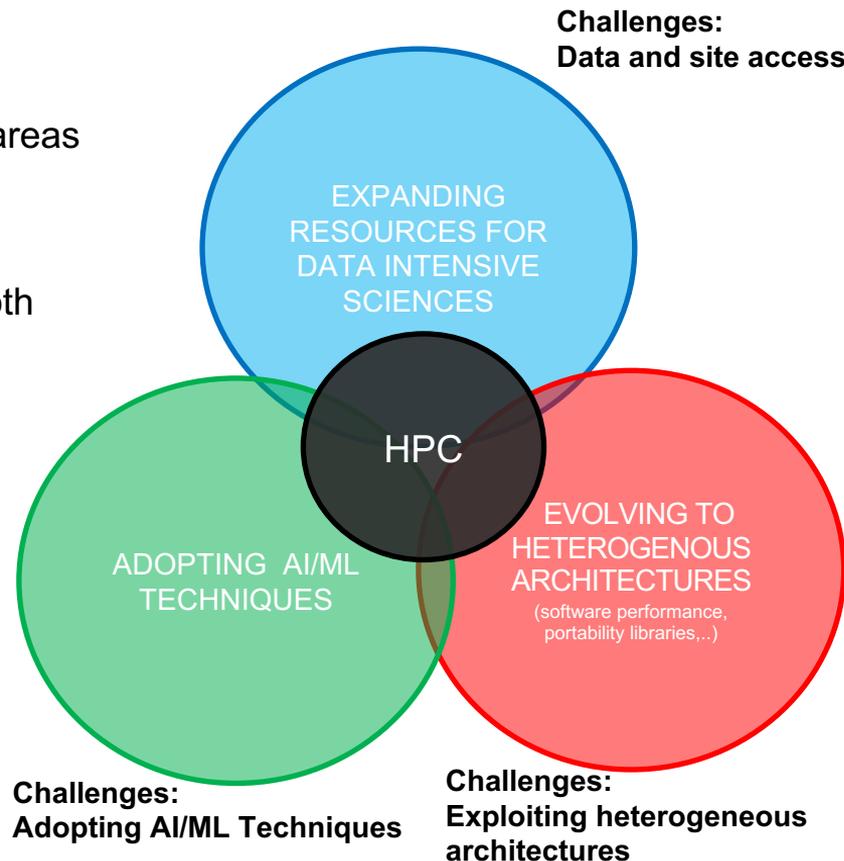
- Expand available computing resources
- Evolve to heterogenous architectures
- Adopt new techniques like AI and ML

HPC integration challenges were compiled by the experiments

- [Common Challenges Document](#)

Recently formed an HPC collaboration between PRACE, GÉANT, SKA and CERN on:

- Data Access
- Authentication and Authorization
- Benchmarking
- Training and Center of Expertise

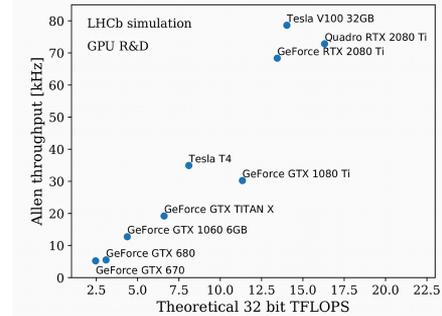


[From the HPC Collaboration Kick-off-Workshop](#) and [Signature Ceremony](#)

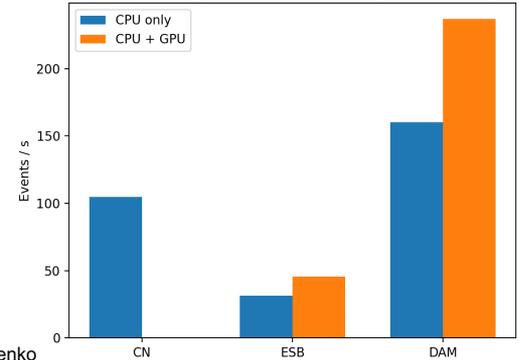


# Exploiting Heterogenous Architectures

- Explorations in heterogeneous hardware are one of the drivers of innovation in computing
  - Large improvements recently in processing performance have come from offloading work to GPUs. Some examples:
    - **CMS High Level Trigger reconstruction** with GPUs increased throughput by ~30% ([Patatrack](#))
      - work done in collaboration with the DEEP-EST project
    - **The LHCb Allen Project** is a high level trigger implemented fully on GPUs ([Allen Project](#))
- Reengineering the software is necessary
  - Results in performance gains, easier adoption of heterogeneous hardware and better maintainability
  - Our codebase is the work of thousands of developers over many years

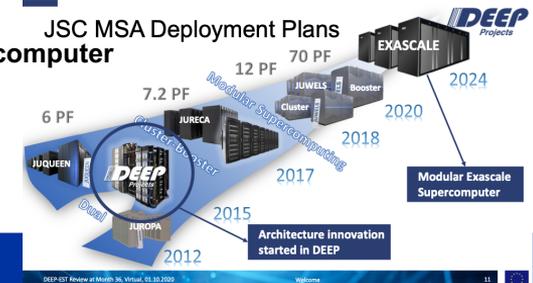


Throughput by node type. CMS HLT Run3 configuration with Open Data



V. Khristenko

## JSC MSA Deployment Plans Modular Supercomputer Architecture



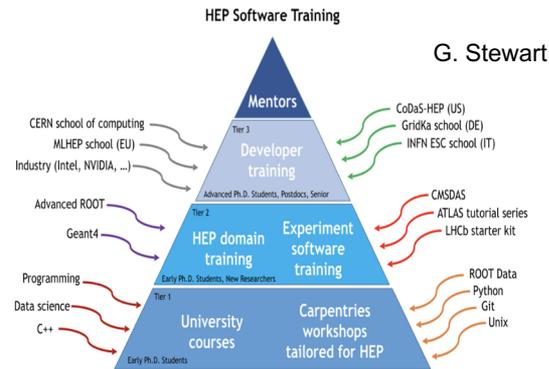
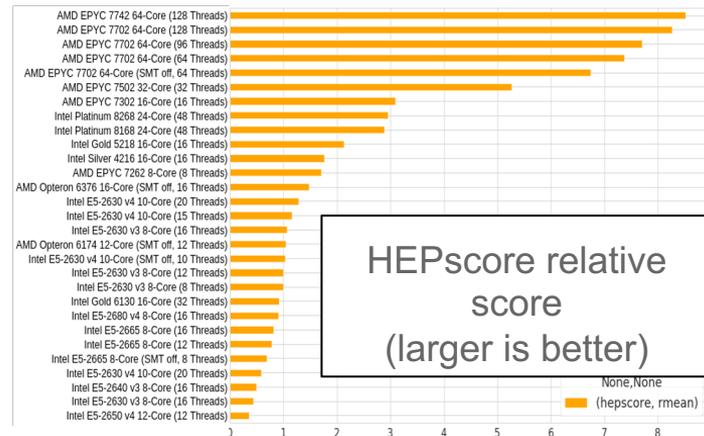
# Exploiting Heterogeneous Architectures

**Evaluate Progress:** Benchmark applications to show we can effectively use HPC systems

- Standalone containerized benchmark suite has been developed
  - Automated collection & reporting of capabilities across several Intel/AMD + Nvidia processors
  - Build on collaboration expertise including augmenting the Unified European Applications Benchmark Suite with HEP applications

**Training:** Build a common center of expertise

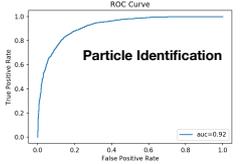
- Make training relevant, scalable and sustainable in HEP
- Dedicated programs throughout 2021 on accelerator programming and performance tuning



# Deploying AI on HEP workflows

The LHC experiments are working on machine learning applications across the whole data collection, processing and analysis

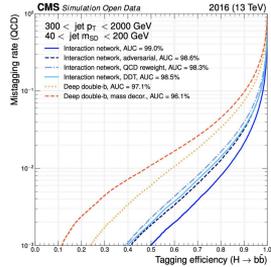
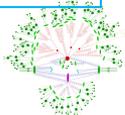
Filtering and Classifying



CERN openlab project with Micron

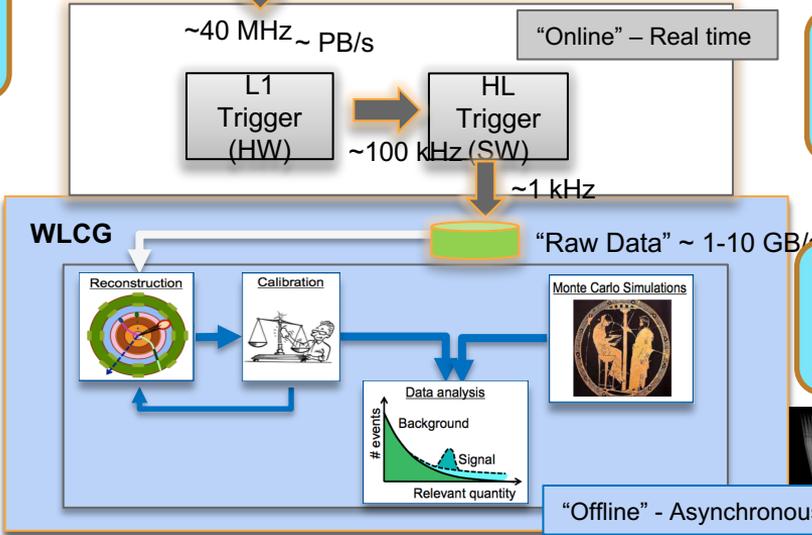
Reconstruction

Opportunities for collaboration with industry, other sciences and the HPC community



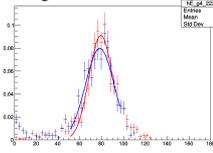
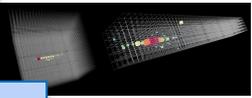
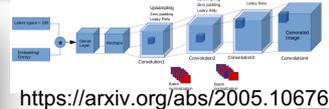
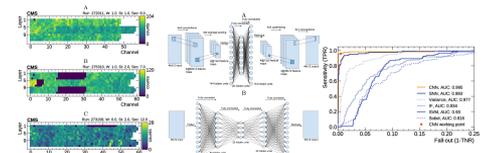
<https://arxiv.org/pdf/2003.11603.pdf>

Analysis



Data Quality Monitoring

Fast Simulation



# Distributed Training for ML workflows

MPI Based distributed training across GPU nodes

In order to use HPC centers for ML distributed training techniques are necessary

- Many parallel hardware elements that are well suited to machine learning training can be used

Opens the possibility for **training and inference as a service** at HPC centers

- Large scale resources for faster training and turnaround
- Performance of HPC sites enables possibility for interactive training via tools like Jupyter, but requires flexible data access to use HPC seamlessly (*DataLake*)

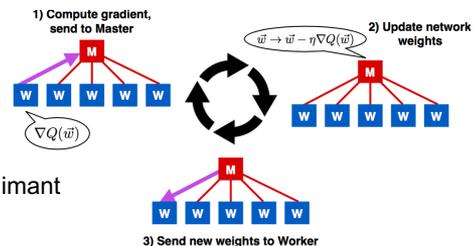
Scaling across multiple CPU nodes

| Nodes | Training Time(S) per Epoch | Linear Time(S) per Epoch | Scaling Efficiency |
|-------|----------------------------|--------------------------|--------------------|
| 4     | 3806                       | 3806                     | -                  |
| 8     | 1910                       | 1903                     | 99.6%              |
| 16    | 1001                       | 951.5                    | 95.1%              |
| 32    | 504                        | 475.75                   | 94.4%              |
| 64    | 253                        | 237.87                   | 94%                |
| 128   | 124                        | 118.93                   | 95.9%              |
| 256   | 61                         | 59.46                    | 97.5%              |
| 512   | 33                         | 29.73                    | 90.1%              |

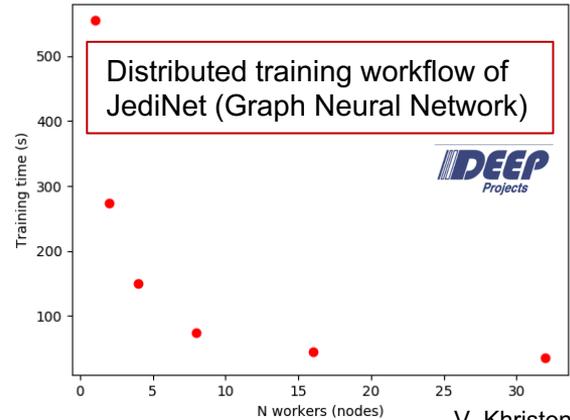
Shows the time taken in seconds for a single epoch of the CERN 3D-GAN network to complete.

<https://arxiv.org/abs/2005.10676>

S. Vallecorsa



J.-R. Vlimant



V. Khristenko

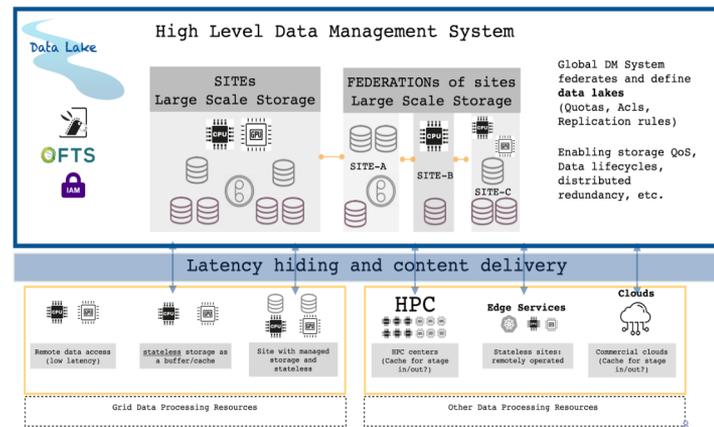
Original work CHEP2019:  
[https://indico.cern.ch/event/773049/contributions/3474799/attachments/1937869/3212749/NNLO\\_-\\_CHEP2019.pdf](https://indico.cern.ch/event/773049/contributions/3474799/attachments/1937869/3212749/NNLO_-_CHEP2019.pdf)

<https://arxiv.org/abs/1811.04492>  
<https://arxiv.org/abs/2009.04509>

# Challenges in Data and Site Access

Using supercomputers for data intensive science will require understanding how to access data efficiently

- Demonstrate a sustained input data rate
  - For HL-LHC a large HPC could process 10PB/day or ~1Tb/s
  - Goal for 2021: multi 100Gb/s
- Stress-test dynamic delivery of data using an HPC
  - HEP is developing a *DataLake* concept to improve data access across resources
- Use of caching layer, potentially as an edge service, and content delivery
  - proof-of-concept early next year



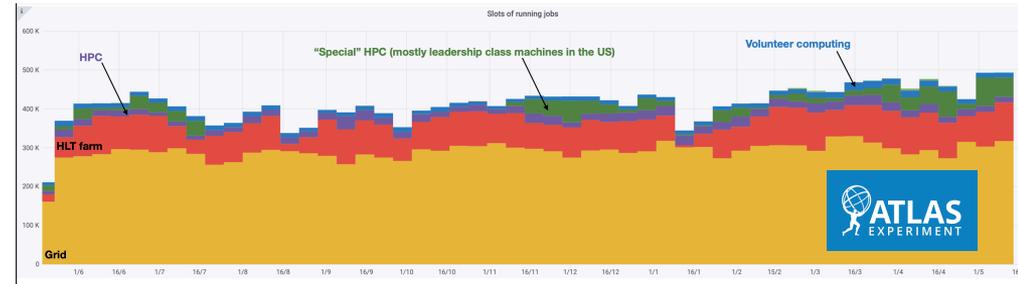
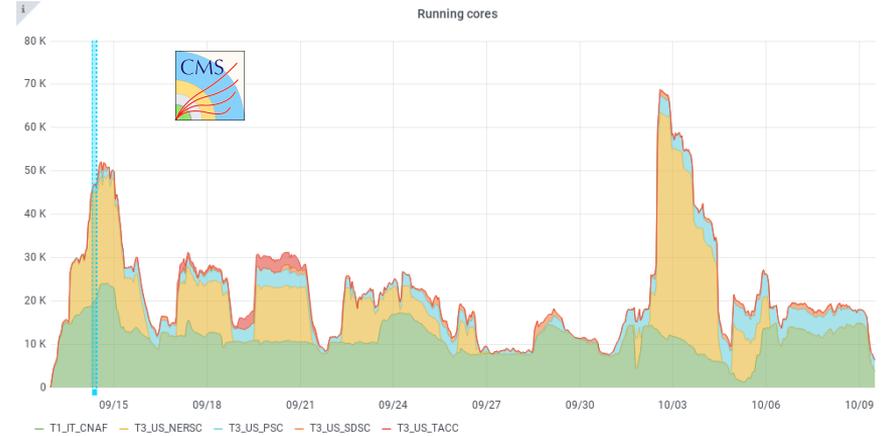
# Accessing HPC Resources

The LHC experiments have demonstrated the ability to exploit individual HPC facilities

- Primarily x86 based systems
- CMS shows usage on 5 sites
- ATLAS plot shows the HPC resource contribution to the total

Access to some PRACE Tier-0 systems for testbed use through the HPC collaboration in 2021

- Will improve access to heterogenous architectures at scale



## PRACE | Tier-0 Systems in 2020

|                                                                                                                                                                                                              |                                                                                                                                                             |                                                                                                                                                                                                |
|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
|  <p><b>ManoNostrum:</b> IBM BSC, Barcelona, Spain #18 Top 500</p>                                                         |  <p><b>Pig Delle:</b> Cray XC50 CDS, Lugano, Switzerland #10 Top 500</p> |  <p><b>NEW ENTRY 2018 JUWELS (Module 1):</b> Atos/Bull Sequana GAUSS @ T3Z, Jülich, Germany #39 Top 500</p> |
|  <p><b>NEW ENTRY 2018 JOLIOT CURIE:</b> Atos/Bull Sequana X1000, GENCI @ CEA, Bruyères-le-Châtel, France #14 Top 500</p> |  <p><b>MARCONI-100:</b> IBM CINECA, Bologna, Italy #9 Top 500</p>       |  <p><b>NEW ENTRY 2019 SuperMUC NE:</b> Lenovo cluster GAUSS @ LZJ, Garching, Germany #13 Top 500</p>       |
|  <p><b>NEW ENTRY 2020 HAWC:</b> HPE Apollo GAUSS @ HRS, Stuttgart, Germany</p>                                          |                                                                                                                                                             |                                                                                                                                                                                                |

Close to 110 Petaflops total peak performance

# Outlook

- By HL-LHC HEP will need exascale computing and storage resources and HPC facilities will play a critical role
  - Expanding our computing resources
  - Expanding our use of heterogenous architectures and access to testbeds
  - Facilitating the use of AI and ML on HPC systems
- We are engaging with the HPC community to tackle the integration challenges
  - Connection to European HPC systems will be very important to prove technical benefits
    - HPC collaboration with PRACE, GÉANT, CERN, and SKA is an excellent start and we will provide application knowledge for a joint program on software optimisation on HPC systems
    - We would like to strengthen further ties with EuroHPC