CPAN
Ingenio 2010

HPC use in WLCG and the Spanish case - J. Flix

XV CPAN days 2023

Santander (Spain)

02-06 October 2023

PIC
port d'informació
científica

IFAE Ciemat RES RED ESPAÑOLA DE España | digital 2026 Unió Europea Fons Europeu Next Generation GOBIERNO DE ESPAÑA Plan de Recuperación, Transformación y Resiliencia Next Generation Catalunya Generalitat de Catalunya Departament de Recerca i Universitats GOBIERNO DE ESPAÑA VICEPRESIDENCIA PRIMERA DEL GOBIERNO MINISTERIO DE ASUNTOS ECONÓMICOS Y TRANSFORMACIÓN DIGITAL SECRETARÍA DE ESTADO DE DIGITALIZACIÓN E INTELIGENCIA ARTIFICIAL

Barriers to exploiting HPC

PIC port d'informació científica

Prospects for LHC computing in Spain for the next 3 years - J. M. Hernández

Exploiting supercomputers for LHC

- Data intensive computing with HPC facilities is a **challenge**
  - Limited/no network connectivity in compute nodes
  - Limited storage for caching input/output event data files
  - **In practice only run cpu-bound workflows (MC simulation) with little I/O**
- Our applications are not really suited for HPC
  - No large parallelization (no use of fast node interconnects)
  - No substantial use of accelerators (GPU)
- Substantial **integration** work to make HPC work for HTC
  - No one-fit-all solution: each facility is different
  - Little effort available in the LHC experiments; in charge of the local communities
  - Experiments do not accept capacity from HPCs as pledged resources unless they can be used transparently as any other WLCG site
- Not suitable resource **allocation** model
  - We would need a guaranteed share of resources rather than apply for allocations

Then… why?

A great **rapidly growing resource** (more than 100x WLCG; WLCG estimated at ~30 PFlops…)

Potential opportunities for **allocations** or **“free” opportunistic** computing usage: can help cover increasing CPU needs despite flat funding

WLCG computing is done economically on the sort of hardware used on the Grid. **National computing priorities** may intend to **complement HTC pledges with HPC resources** at some point soon

Interesting **R&D:** access and use of heterogeneous resources (GPUs, ARMs, POWER)

Various barriers for exploitation have been reduced over the years, **some still exist…**

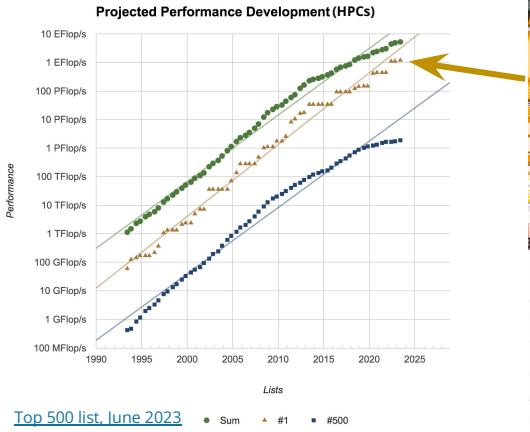

Top 500 list, June 2023

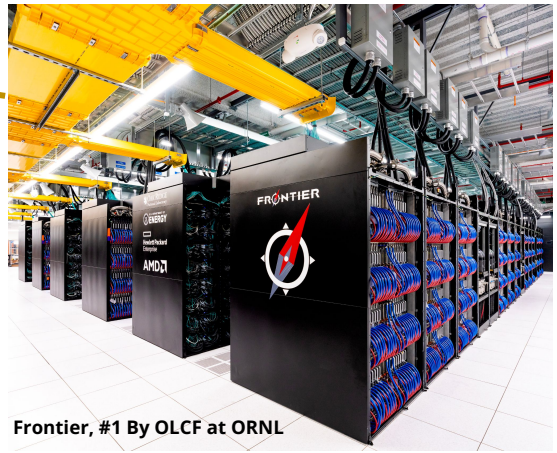

Frontier, #1 By OLCF at ORNL

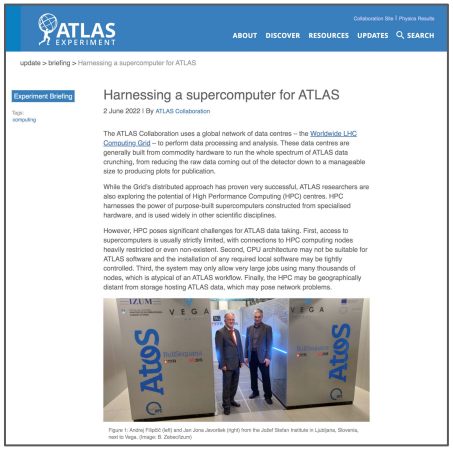

ATLAS briefing on Vega HPC, June 2022

HPCs used per experiment

Usage of HPCs in WLCG at the end of 2022

- LHCb
  - Piz Daint in CSCS (Switzerland)
  - Marconi-A2 in CINECA (Italy) – not used anymore
  - SDumont in LNCC (Brazil)
  - MareNostrum in BSC (Spain)
- ALICE
  - CINECA (Italy) - used in 2020
  - LBNL (Berkeley, USA):
    - Lawrencium (in production)
    - CORI (to be decommissioned)
    - Perlmutter (being commissioned)
- ATLAS
- Past:
  - Titan (Oak Ridge), Theta (Argonne), Edison (NERSC), Tianhe-1 (China), SuperMUC (Germany)
- Currently:
  - Frontera (TACC), Cori (NERSC), Vega and Karolina (EuroHPC), Toubkal (Morocco), RIVR (Slovenia), Tokyo HPC (Japan)
  - As part of pledge: MareNostrum (BSC), CSCS (CH), HPC2N and NSC (NDGF)
- Planned:
  - Perlmutter (NERSC) currently being commissioned

CMS:

Machine	Location	Architecture*	Status
Piz Daint	CH (CSCS)	x86 + Nvidia GPU	Production
CLAIX	DE (RWTH Aachen)	x86 + Nvidia GPU	Production
HoreKa	DE (KIT)	x86 + Nvidia GPU	Production
Marconi 100	IT (Cineca)	POWER9 + nVidia GPU	Validated, pre-production
MareNostrum 4	ES (BSC)	X86 (+ GPU)	production
Cori	US (NERSC)	x86	Production
Frontera	US (TACC)	x86	Production
Stampede2	US (TACC)	x86	Production
Bridges-2	US (PSC)	x86	Production
Expanse	US (SDSC)	x86	Production
Anvil	US (Purdue)	x86	Production
Perlmutter	US (NERSC)	X86 + nVidia GPU	integration/commissioning
Summit	US (OLCF)	Power9 + nVidia GPU	integration/commissioning
Frontier	US (OLCF)	X86 + AMD GPU	Planned
Polaris	US (ALCF)	X86 + nVidia GPU	Planned
Ookami HPC testbed	US	ARM	Planned
Leonardo	IT (CINECA)	X86+ nVidia GPU	Planned
MareNostrum 5	ES (BSC)	x86 / Arm + nVidia GPU (tbc)	Planned
Jureca	DE (FZJ)	X86 + nVidia GPU	Planned

CMS HPC use in M core hours per month Jan 2020 - Aug 2022

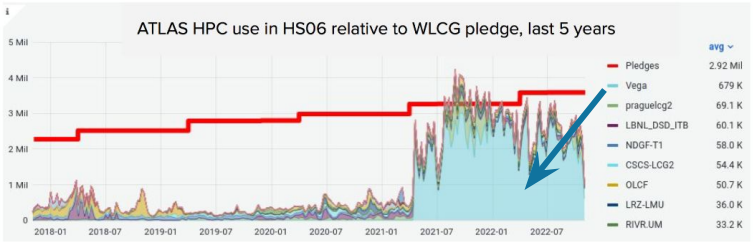

David Cameron, HPC at WLCG workshop, Lancaster 8.11.22

Suitable workflows in HPCs

The main workflow is **MC generation/simulation elsewhere**

- Easy on I/O, small or no input data, low memory and few dependencies on external services
- Stable software releases used for long periods
- Not time-critical

Where possible, **other production workflows are run** (eg CMS full chain on all DOE/NSF HPC)

Some however run **all workflows including analysis**

- Vega, CSCS and NDGF Tier-1 for ATLAS
- HPCs at LBNL for ALICE (helped by co-located T2 site)
- CSCS, CLAIX and HoreKa for CMS

Friendly to HEP HPCs!

ATLAS activities executed at CSCS

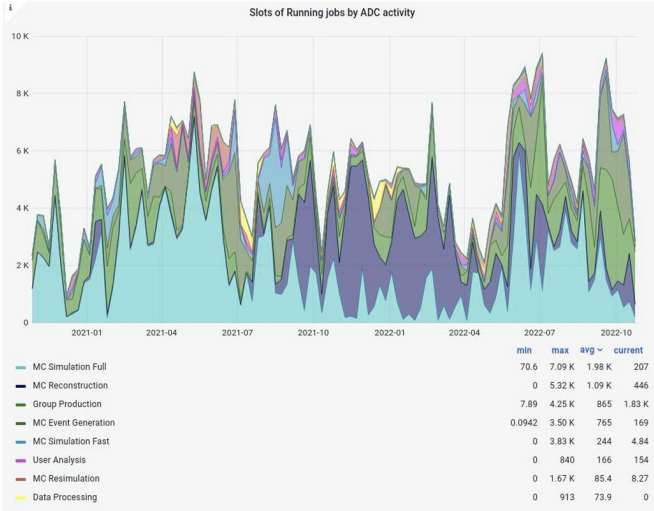

Software availability

PIC
port d'informació
científica

CVMFS is the de facto standard throughout HEP for **software distribution**

Some HPCs install CVMFS + squids natively, which makes things a lot easier

Various **solutions exist for HPC without CVMFS**

- Copy part of the tree (eg particular software release) or the whole tree (with **cvmfs_preload**) to local file system and point jobs there
- Install CVMFS as an unprivileged user with **CVMFSExec**
- Develop tools such as **subcvmfs-builder** to deploy releases automatically
- Pack software into a **container** in which the job runs

Some of these solutions are generally only suitable for HPC running **specific releases** of a **specific workflow**

To make use of HPC as a general purpose resource **CVMFS must be available natively**, or **at least a squid server** provided to make use of CVMFSExec

- Copying the whole tree (with **cvmfs_preload**) to local HPC file system, and keeping updates, also works, indeed

Data management and communication

Two **main issues to solve**:

- Local storage, but no local storage element
- No network connectivity to access remote storage or experiment services

Some sites provide **proxies** or **gateways** for network traffic

Others: **“tunnelling” traffic** through a login or edge node, or traffic through shared file-systems

Typically, this is solved by deploying an **edge service** which handles data transfers (if possible!)

How does the **pilot model wor**k in the most heavily constrained environments?

- Pilot runs on the edge node or even outside the HPC: pull job, do data transfer then submit payload to batch system → **DIRAC PushJobAgent**
- Edge service which handles job communication, data transfers and submits to batch → **ATLAS Harvester**
- Central service which handles job communication, and submits fully-formed jobs (payload and input/output) to edge service which handles data transfers → **ATLAS ARC Control Tower** (central) and **ARC-CE** (at HPC or HTC center)
- Edge service which handles communication with jobs through shared file system: e.g. **CMS at BSC**

Other issues to solve

Access to external services for e.g. **conditions data (Frontier/Squid)**

- “Fat” container image with all software and conditions data
- Mirror to local shared file system
- Edge service proxy
- Use “tunnels” through a login or edge node

Batch system policies

- Whole node scheduling is required
  - Single pilot running many parallel single core jobs inside
  - Multi-threaded/multi-core jobs using whole node
- Minimum nodes per job
  - “Fat pilots” running many jobs in a single batch job, ATLAS “Jumbo jobs”, CMS Glideins with partitional slots, …

In general, **significant effort** is required from central experiment teams and site contacts (the people who have access to HPC facilities and submit allocation proposals)

HPC lifespan is short (3-5 years) compared to WLCG sites

- Continuous work to commission the next new one (although can be made easier by existing relationships with sites)

Integration of BSC CPU resources

In 2020 BSC designated LHC computing as a **strategic project**

- Agreement promoted by WLCG-ES community and funding agency

Allocations* of up to a **7% share of MN4 for LHC** [100M coreHours max/year]

- ~70M coreHours/year

* Submission of proposals for CPU time allocation every 4 months

Potentially, very **significant contribution** for LHC computing in Spain

- Comparable e.g to all ATLAS+CMS+LHCb simulation needs in the country

Integration of BSC CPU resources

PIC
port d'informació
científica

BSC - Barcelona Supercomputing Center

- Largest HPC center in Spain
- Current MareNostrum 4 (*MN4*) general-purpose cluster:
  - 11.5 Petaflops (166k CPU cores), 390 TB RAM, 24 PB disk local SSD disk
  - SLURM as batch system, SUSE Linux Enterprise as OS
  - 15 PB GPFS as storage back-end (mounted on login/compute nodes)
- Next MareNostrum5 (~17xMN4, ~200 petaflops)
  → being commissioned and available soon!
  - One of Europe's first pre-exascale supercomputers

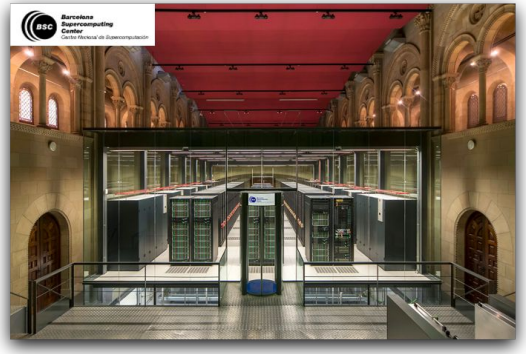

https://eurohpc-ju.europa.eu/about/our-supercomputers_en#marenostrum-5

BSC imposes very **restrictive network connectivity** conditions

- No incoming *or outgoing* connectivity from compute nodes
- Only incoming SSH/SSHFS communication through login nodes
- A shared disk (GPFS) mounted on compute nodes and login machines - accessible from outside via sshfs
- No services can be deployed on edge/privileged nodes

Use of the BSC resources

Services installed at the Spanish WLCG sites to access and exploit BSC resources:

- **PIC Tier-1**: **2x ARC-CEs** for both **ATLAS** and **LHCb**, and **custom-made gateway** for **CMS**
- **IFIC Tier-2**: **1x ARC-CE** for **ATLAS**
- **UAM Tier-2**: **1x ARC-CE** for **ATLAS**

Submission of **proposals for time allocation** every 4 months

- 3x proposals for **ATLAS** (A. Pacheco-IFAE, S. González-IFIC, J. del Peso-UAM)
- 1x proposal for **CMS** (J. Flix-CIEMAT)
- 1x proposal for **LHCb** (X. Vilasis-Ramon Llull)

Since 2020, **57 allocation proposals have been submitted** and approved by BSC

- Allocations sum up **161 million CPU hours**
- We have utilized **157 million CPU hours**
  - The current allocations overall usage is at 75% usage → +5 million CPU hours until end of October 2023

Taking into account the period length and CPU power at BSC, this utilization corresponds to an average installed capacity of approximately 82 kHS06, representing around **40% of the average Grid resources deployed in Spain** for the LHC experiments in the period 2020-2023

Use of the BSC resources

ATLAS

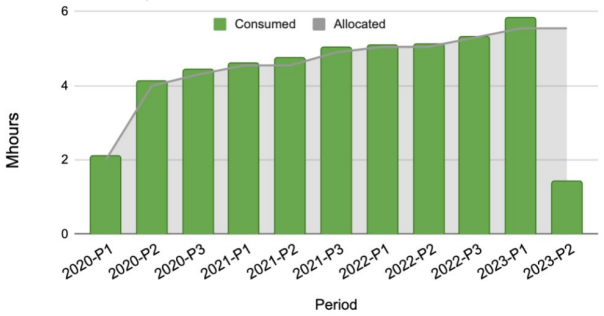

[PIC ARC-CE gateway]

CMS

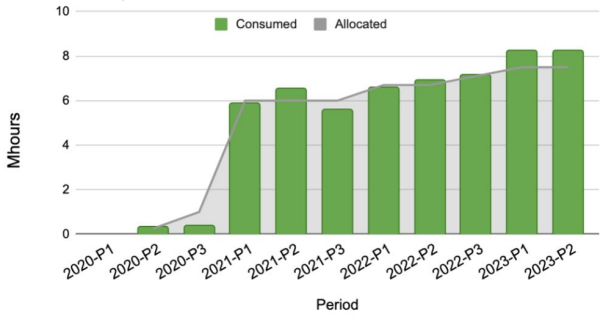

[PIC custom-made gateway]

LHCb

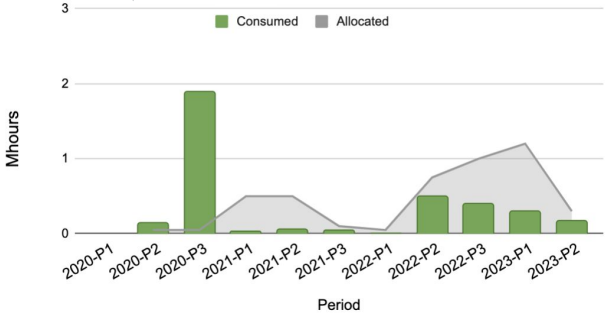

[PIC ARC-CE gateway]

ATLAS

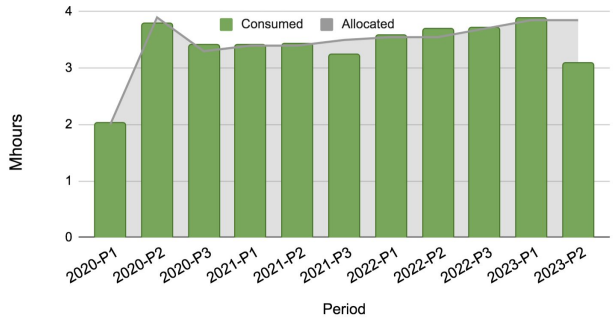

[IFIC ARC-CE gateway]

ATLAS

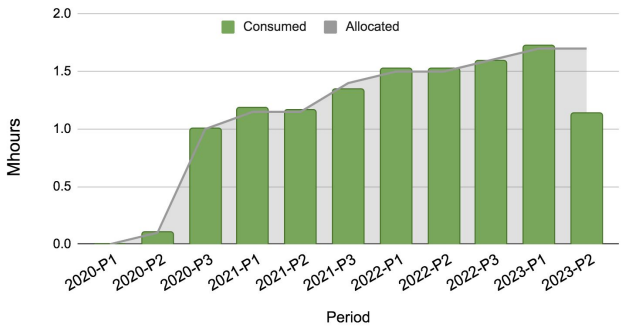

[UAM ARC-CE gateway]

Use of the BSC by ATLAS

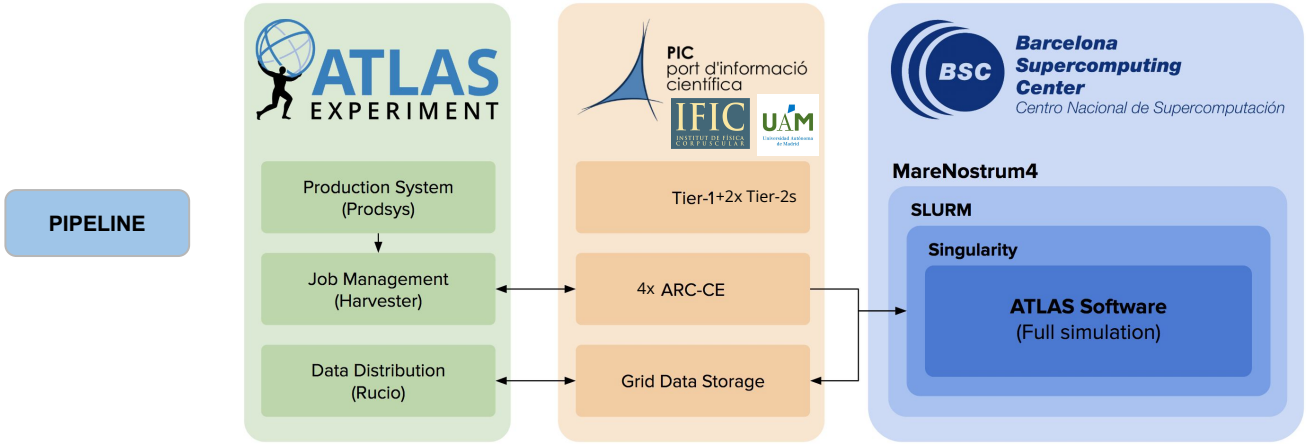

Submitting **ATLAS** payloads to BSC since 2018, in production since 2019

Using four **ARC-CEs** in Spain to interconnect MareNostrum and ATLAS production system

Only simulation workflow validated - singularity containers, pre-placed at MareNostrum GPFS

~33 million CPU hours used at BSC **last year** by ATLAS through these gateways

→ **At CHEP2021 proceedings (link)**
→ **At CHEP2023 (link)**

Use of the BSC by LHCb

LHCb used similar technical implementations as ATLAS (**ARC-CE02.PIC.ES**) to exploit BSC resources - submitting grants to BSC as ATLAS and CMS, and **modified DIRAC** for the purpose

~2.5 million CPU hours used at BSC **last year** by LHCb through this gateway

LHCb

Node ES-PIC — Total number of jobs by Submit Host and Month (Custom VOs)

Submit Host	Oct 2022	Nov 2022	Dec 2022	Jan 2023	Feb 2023	Mar 2023	Apr 2023	May 2023	Jun 2023	Jul 2023	Aug 2023	Sep 2023	Total
ce13.pic.es:9619/ce13.pic.es-condor	14,744	14,141	12,499	24,009	14,431	14,351	18,780	27,443	34,227	50,111	34,072	41,242	300,050
ce14.pic.es:9619/ce14.pic.es-condor	13,371	14,006	12,116	24,172	13,654	17,186	17,397	28,819	35,737	49,131	29,967	34,655	290,211
gsiftp://arc-ce02.pic.es:2811/jobs	0	0	0	0	0	0	4	2	0	0	0	0	6
https://arc-ce02.pic.es:8443/arex	16,836	463	15,714	21,962	9,387	16,759	6,338	10,727	10,923	42,703	15,434	14,452	181,698
Total	44,951	28,610	40,329	70,143	37,472	48,296	42,519	66,991	80,887	141,945	79,473	90,349	771,965
Percent	5.82%	3.71%	5.22%	9.09%	4.85%	6.26%	5.51%	8.68%	10.48%	18.39%	10.29%	11.70%	

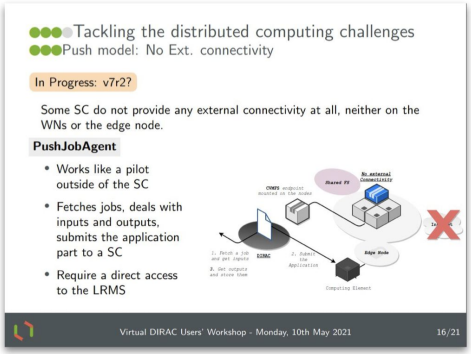

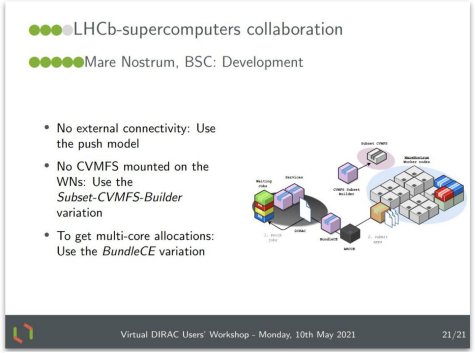

A.Boyer, "Integrating DIRAC workflows in Supercomputers" [link]

Use of the BSC by CMS

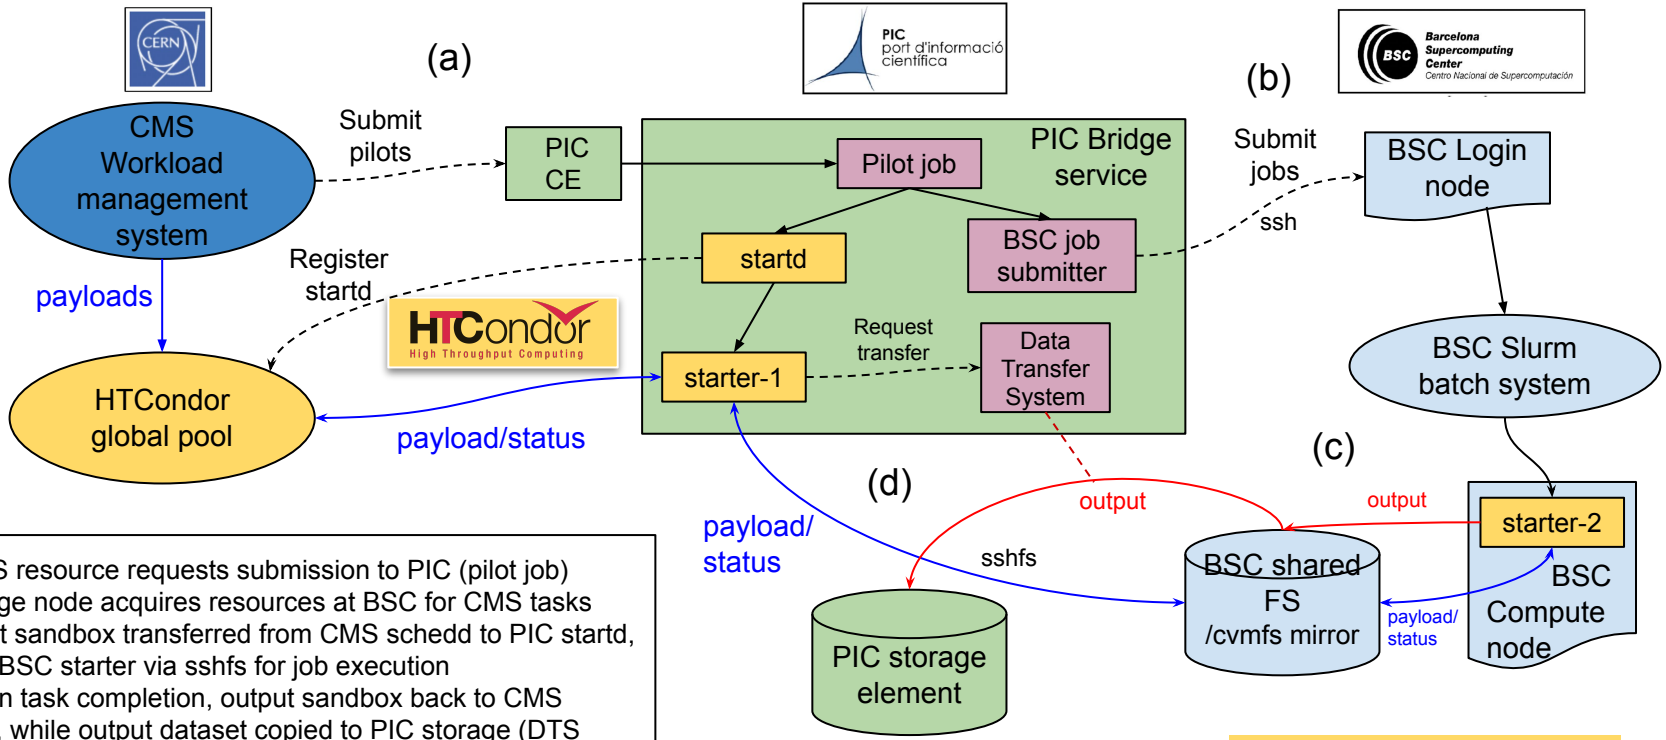

(a) CMS resource requests submission to PIC (pilot job)
(b) Bridge node acquires resources at BSC for CMS tasks
(c) Input sandbox transferred from CMS schedd to PIC startd, then to BSC starter via sshfs for job execution
(d) Upon task completion, output sandbox back to CMS schedd, while output dataset copied to PIC storage (DTS acting as third party copy manager)

→ **At CHEP2021 proceedings (link)**
→ **At ISGC 2022 (link)**
→ **At CHEP2023 (link)**

PIC and HTCondor team collaboration to **use a shared FS as control path for HTCondor**

Use of the BSC by CMS

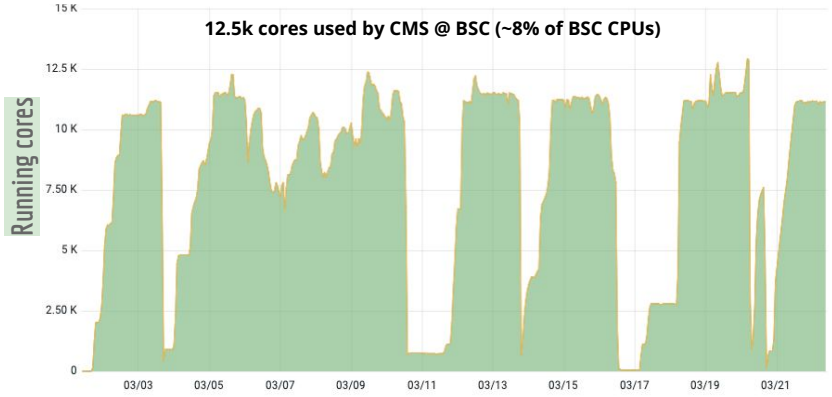

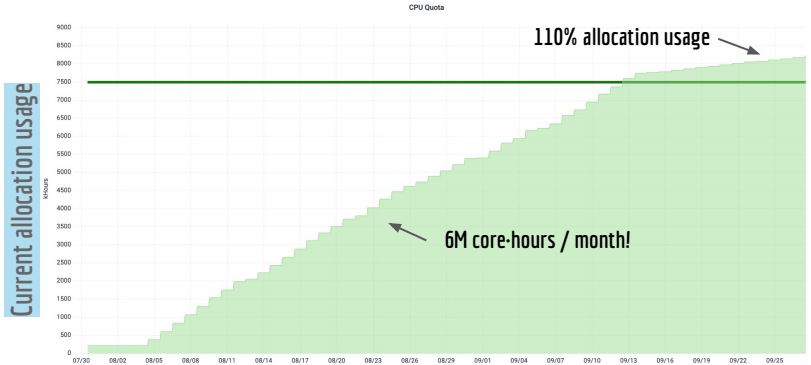

Running in operations (**CMS Simulation workflows GEN-SIM**)

Result of the **PIC and HTCondor team collaboration** to use a **shared FS as control path for HTCondor**

Interaction with BSC execute nodes through the login node, mounting the shared FS through **sshfs** and sending jobs to the Slurm scheduler via **ssh.** Slurm jobs launch a HTCondor slot that joins the **CMS Global Pool**

Using **cvmfs_preload** to bring cvmfs CMS files to BSC. Two weeks to copy ~37M files (13 TB), at first injection. **cvmfsexec** used to build the cvmfs file structure

Stage-in/out + **Data Transfer Manager** designed to transfer input and output data from/to PIC (*xRootD server in singularity images*)

Integrated with **WMAgent @ CERN** + **Accounting to APEL**

~22 million CPU hours used at BSC **last year** by CMS through this gateway

→ **At CHEP2021 proceedings (link)**
→ **At ISGC 2022 (link)**
→ **At HTCondor WS 2022 (link)**
→ **At CHEP2023 (link)**

Next steps

Still, **further refinements can be applied** to our setup to improve automated operations and quality of the service

- Enabling and testing all these functionalities in the **upcoming MN5 facility** [link]
- Enhanced **network connectivity** between BSC and Spanish WLCG sites
  - BSC expects 100 Gbps along 2023 (current 2x10 Gbps)
- Improving **monitoring**, **alarms** and **automated recovery** for the services
- **Optimize resource usage efficiency** with closer connection to experiments
- Possible use of the systems in place to **manage also input datasets** would increase the usability of the resources by the experiments
  - i.e. data reprocessing and analysis tasks
- Potential for **GPU and ARM resources** exploitation through MN5 facility

Spanish WLCG Pledges

Most **HPC use in WLCG is opportunistic**, i.e. outside the WLCG pledge framework

Historically, only if HPC resources **look like a grid site**, one can pledge resources to WLCG

- Already done for CSCS and parts of NDGF-T1
- INFN will pledge Leonardo (CINECA EuroHPC) as part of Italian T1 resources from 2023

In order to reduce Grid project costs, the **CPU contribution of BSC will be part of the Spanish WLCG pledge from 2024 on** (even if we are not yet running all type of workflows)

- The plan is to run all of the ATLAS, CMS and LHCb simulations at BSC (~50% of the cpu pledge). Hence we requested to Plan Nacional funding for only 50% of the required CPU capacity at Grid sites

The agreement with BSC ensures that **resources must be constantly available** at some level with a multi-year commitment

- If the agreement is not preserved, there is a risk of losing CPU cycles in the next years

We project to use at BSC **68M hours in 2024**, **92M hours in 2025** and **115M hours in 2026**

Conclusions

A lot of **great work has been done** to exploit HPCs in WLCG. In particular, a difficult context has been addressed to use BSC CPU resources by ATLAS, CMS and LHCb

So far the exploitation since 2020 represents around 40% of the average grid resources deployed in Spain - **not far from the 50% target** for exploitation set for 2024 on

The work done to integrate this resource took **many FTE efforts** from the WLCG Spanish community and from international teams (HTCondor and experiment frameworks)

Current challenges on the software side to **exploit different architectures** (GPU, ARM). Door for opportunities to test/use these type of resources at MN5?

Some HPC facilities look like they are **becoming more friendly to HEP**, at least in terms of accessibility. We do expect the **same level of difficulties in MN5**

Since HPC facilities are constantly being designed… Can WLCG present a **united front** and **have a voice** to influence?

- Joint ECFA-NuPECC-APPEC (JENA) workshop resulting in a Working Group to focus on this, including Particle Physics, Nuclear Physics, and Astroparticle Physics communities (data intensive sciences)

Thank you!

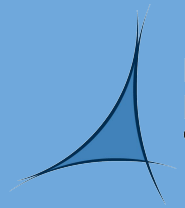

PIC
port d'informació
científica