

# Enabling ATLAS big data processing on Piz Daint at CSCS

**F G Sciacca on behalf of the ATLAS Collaboration**

Albert Einstein Center for fundamental Physics, University of Bern, Sidlerstrasse 5, CH-3012 Bern, Switzerland

E-mail: [gianfranco.sciacca@lhep.unibe.ch](mailto:gianfranco.sciacca@lhep.unibe.ch)

**Abstract.** Predictions for requirements for the LHC computing for Run 3 and Run 4 (HL-LHC) over the course of the next 10 years show a considerable gap between required and available resources, assuming budgets will globally remain flat at best. This will require some radical changes to the computing models for the data processing of the LHC experiments. Concentrating computational resources in fewer larger and more efficient centres should increase the cost-efficiency of the operation and, thus, of the data processing. Large scale general purpose HPC centres could play a crucial role in such a model. We report on the technical challenges and solutions adopted to enable the processing of the ATLAS experiment data on the European flagship HPC Piz Daint at CSCS, now acting as a pledged WLCG Tier-2 centre. As the transition of the Tier-2 from classic to HPC resources has been finalised, we also report on performance figures over two years of production running and on efforts for a deeper integration of the HPC resource within the ATLAS computing framework at different tiers.

## 1. Introduction

In view of the challenges posed by the foreseen scale of the High-Luminosity LHC runs (2025-2034) computing needs [1], significant R&D efforts have been put in place as part of the upgrade programs, in order to address the predicted shortfall of resources for reconstruction and offline computing resulting from higher trigger rates, larger event size and greater complexity. The current models assume global budgets to be flat and have to assume technology trends that are hard to predict with certainty for the next decade. The gap between required and available resources is currently estimated to be about one order of magnitude [2]. This will require optimisation of several aspects of experiment software, resource provisioning and usage, system performance and efficiency and data processing.

In this context, the integration of general purpose HPC machines with the LHC experiment frameworks has been a hot topic for several years as, in some circumstances, they hold the potential to offer a more cost-effective data processing infrastructure compared to dedicated resources in the form of commodity clusters. However, while the latter can in general be tuned exactly for the type of workloads expected, the former are known to pose higher integration challenges, decreased levels of flexibility and have not generally been operated continuously for LHC experiments workload processing.

The Swiss HEP computing community and CSCS started working on the HPC integration with the LHC experiment Tier-2 facilities in 2014. A Tier-2 facility is currently operated as pledged resource for ATLAS, CMS and LHCb on the *Piz Daint* [3] supercomputer.

## 2. Piz Daint integration timeline

Since the beginning of WLCG operations, CSCS has hosted a WLCG Tier-2 centre on behalf of the Swiss Institute of Particle Physics (CHIPP) in the form of a dedicated commodity cluster (*Phoenix*).

In 2014, the Swiss ATLAS community launched a project aimed at integrating a *Cray XK 7* at CSCS with the ATLAS computing framework PanDA [4]. The integration was accomplished with a non-invasive model: remote submission via ssh using a modified ARC Compute Element hosted at University of Bern [5]. Following the integration phase and some studies about possible performance improvements by recompiling the software with Cray specific flags, we ran ATLAS Geant4 [6] simulations on this system for over six months in production mode. Since the optimisation studies proved inconclusive, we ran the ATLAS unmodified software from CVMFS [7].

In 2015 CSCS and the Swiss HEP community (ATLAS, CMS, LHCb) launched the *LHConCRAY* project, aimed at integrating the current generation supercomputer at CSCS with the computing frameworks of the three LHC experiments. All of the three experiment workflows were targeted, including user analysis. The first phase consisted of implementing a proof of concept on a development *Cray XC40*.

In 2016 the final procurement of dedicated Tier-2 hardware for the Phoenix cluster took place, and in 2017 we started production on Piz Daint on a small scale (1.6k cores) in order to refine the architecture and perform the tuning needed in the production environment of the three experiments. Most of the integration issues were tackled during this phase and a pricing scheme for the resource was agreed for up to the expected lifetime of Piz Daint (2021). This was set to allow a slightly faster pledged computing growth year-on-year, compared with the growth predicted for the dedicated Phoenix cluster. Opportunistic usage of the facility was ruled out by CSCS.

Since scale-up was obviously a concern, in 2017 we performed a test running ATLAS Geant4 simulation on up to 27k cores for an agreed period of time of roughly 6 hours. The test proceeded smoothly, removing immediate scale-up concerns, at least as far as simulation was concerned. It should be noted that the test environment was exclusive to ATLAS and not shared at the same time with other experiment workloads or other users of the facility.

At the end of 2017, it was decided to phase out the dedicated Phoenix cluster and move all Tier-2 workloads onto Piz Daint. For the next years an increase of 25% computing performance at constant funding was projected. The phasing in of Piz Daint has been performed in stages: 4k cores by April 2017, 9k by April 2018. In April 2019 the Tier-2 facility was fully migrated to Piz Daint with 12k cores in production. These resources are dedicated to the WLCG Tier-2 and shared by ATLAS, CMS and LHCb. More details on the various integration steps can be found in ref. [8] and [9].

In addition to the Tier-2 pledged resource commissioning, in 2018 ATLAS worked together with CSCS on a R&D project aimed at commissioning Tier-0 reconstruction of RAW data on demand on Piz Daint on up to 10k cores. In a second phase, CMS joined the effort with an additional 2k cores.

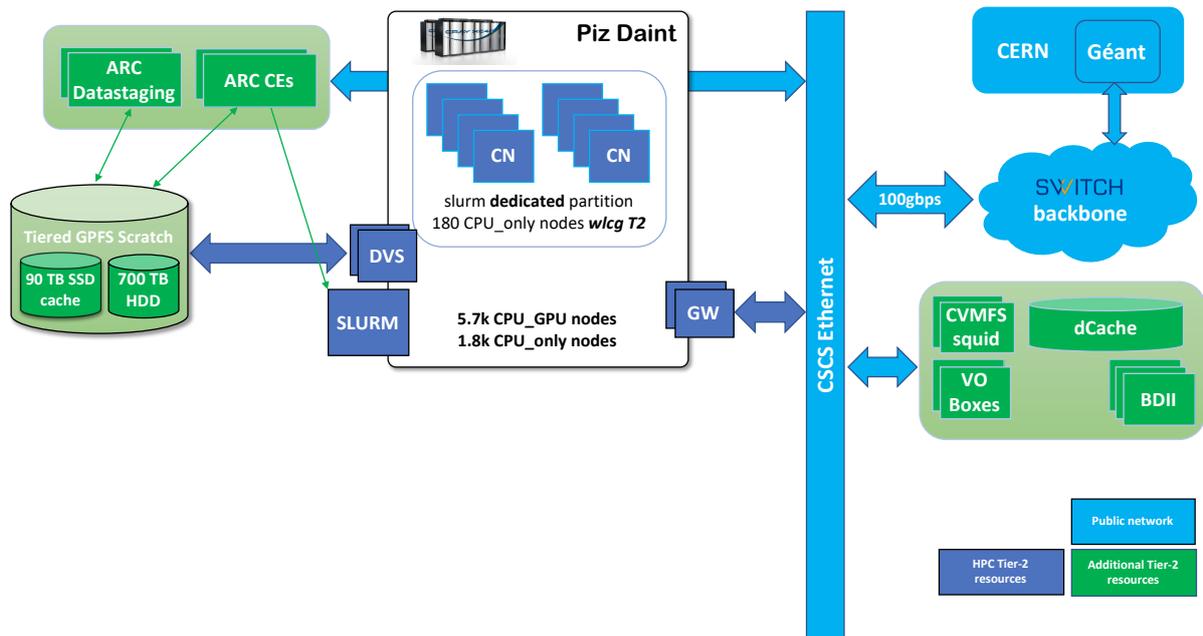
## 3. Tier-2 Integration solutions

*Piz Daint* is a Cray XC50, CPU/GPU hybrid supercomputer, featuring 5704 CPU/GPU nodes and 1813 CPU-only nodes. The full system specifications can be found in [3]. The current solution architecture for the pledged Tier-2 facility is shown in Figure 1.

Classic HPC systems like Piz Daint are typically shared among several communities. They feature cutting edge technologies and user communities architect their software specifically for such systems. With LHC experiments, the challenge is somehow reversed: ATLAS and the other LHC experiments have developed over the years complex and sophisticated distributed

computing frameworks. A large part of the codebase, like event generators and detector simulation toolkits feature legacy code that has historically been developed according to the serial paradigm: events are processed serially on a single thread and embarrassingly parallel processing occurs for scalability purposes in multi-processor systems. HPC systems, on the contrary are usually optimised for scalable parallel software that exploits the tight interconnection between nodes and makes use of accelerators. In addition, network and I/O patterns differ greatly from those of HEP workloads. This raises the necessity of adapting the HPC to such aspects of the HEP computational environment.

On Piz Daint, ATLAS and the other LHC experiments run on a dedicated partition that includes a set of CPU-only nodes. These nodes are reserved for WLCG workloads and not shared with other users in the system. One of the reasons for this is that the WLCG workloads pose the additional requirements of interfacing some services that are external to the Piz Daint integrated network: in particular, mount points are needed for CVMFS and a dedicated GPFS scratch file system. The latter is not one of the two shared Lustre file systems that are integrated with the Cray, and is specifically hardened and tuned for the I/O patterns of WLCG jobs. For a full Tier-2 integration, the classic grid services are also needed: ARC compute elements, dCache storage element, squids for Frontier and CVMFS, information system and VO boxes.



**Figure 1.** CSCS Tier-2 architecture with Piz Daint

### 3.1. Solution architecture highlights

The technical details of the Piz Daint integration are described in great detail in [9]. The decision to use the supercomputer for all the Tier-2 workflow has had an impact on the solutions. Some

of the key points are listed in this section.

- 64 out of the 72 hyperthreading cores on each Piz Daint node are used, since the nodes feature 128 GB of memory. Nodes are set to be *not* job exclusive, single core and multi core jobs are supported.
- Since the nodes are diskless, job execution relies on a shared GPFS file system, featuring an additional SSD based pool acting as cache in order to improve performance both for throughput and *iops*. The GPFS file system acts also as an ARC cache and as a lower tier preloaded cache for CVMFS.
- The Cray *Data Virtualisation Service (DVS)* [10] is used to expose the GPFS file system to the Cray nodes, as such it is a critical link in the integration chain. 19 DVS nodes are currently used to support I/O for 180 Cray nodes. The number is the pragmatic result of operational experience under realistic I/O load.
- CVMFS makes use of a novel two-tier scheme for the cache developed specifically for this application [11]. The upper tier is in-RAM, for which 10 GB per node are allocated for the three experiments. The lower tier is preloaded on the GPFS shared file system.
- The scheduler *SLURM* is configured not to treat the node memory as *consumable*. This allows for memory overcommitment. In addition, jobs are not killed if they exceed the requested or allocated memory quota. Given that memory on the nodes is less than the canonical 2 GB per core (recommended on average for the WLCG workloads), memory handling is another critical element of the integration chain. These settings ensure memory usage peaks from certain workloads can be absorbed due to the natural workload entropy on the nodes. A 6 GB/core limit is enforced in order to catch rogue jobs.
- Job execution is containerised: shifter and singularity are supported. ATLAS has used both, and is currently using singularity only, set as a standard across all sites.

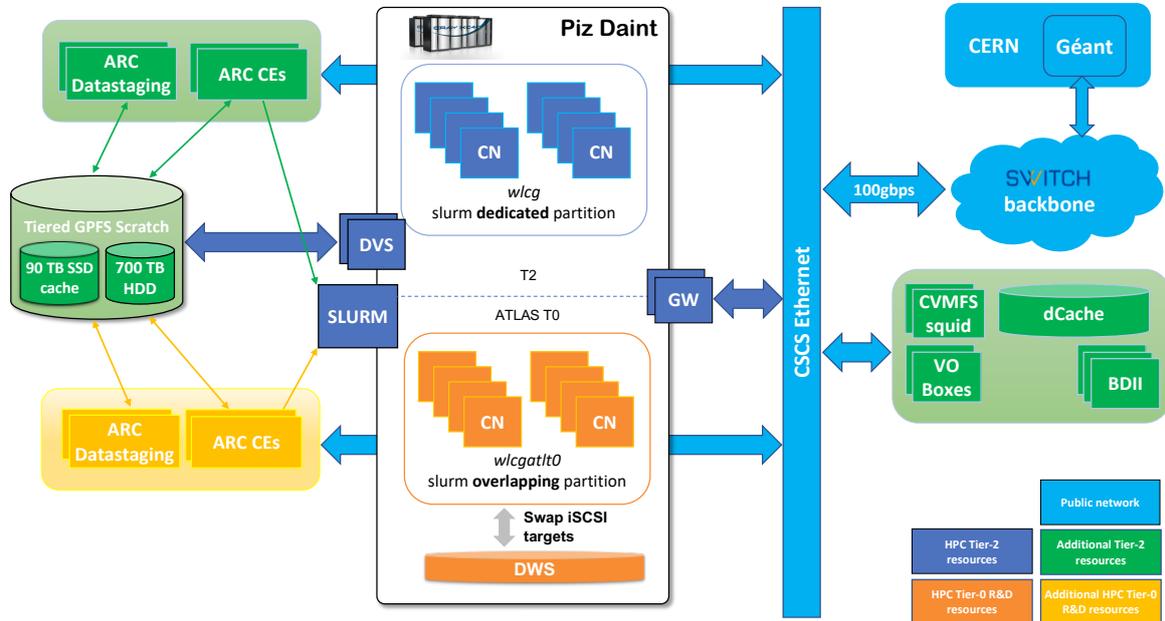
Reaching this level of maturity has been manpower intensive and has required several iterations over the years. The evolving nature of the ATLAS computing model and the increasing demand on resources, as learnt during the production period, also mean that the solution needs constant scrutiny and frequent adjustments. This is challenging, since, as pointed out earlier, the LHC computing demands do not align with general HPC systems. The other consideration to be made is that some of the solutions developed for Piz Daint are system specific and could be invalidated by future machine choices. The extent of this is currently hard to predict.

#### 4. Beyond tier-2 Integration: Tier-0 workloads

A deeper level of integration of the CSCS supercomputer was pursued by commissioning, as a part of an R&D project, the processing of ATLAS Tier-0 workloads on Piz Daint.

ATLAS RAW data reconstruction for each LHC run is performed at the Tier-0 at CERN on dedicated “hardened” resources. This workload is memory hungry and I/O intensive, exceeding the Piz Daint specifications. Whenever a backlog of data to process builds up at the Tier-0 (depending on the LHC performance), ATLAS might want to offload (or *spill-over*) some runs to other sites. Spill-over distributed to the grid was commissioned in 2018 by ATLAS, and in parallel, the possibility to spill-over “on demand” to a single centre as CSCS was explored.

The solution architected for CSCS is based on the Tier-2 architecture, with one crucial difference: nodes used for the processing are no longer reserved and therefore permanently customised for WLCG workloads. They are instead added to a SLURM partition that overlaps with other general users of Piz Daint. When ATLAS Tier-0 workloads are sent to CSCS, they are assigned a much higher priority in the scheduler by means of a *Quality of Service (QoS)* policy. They are then able to run as soon as one or more nodes are free from other user jobs.



**Figure 2.** CSCS Tier-2 architecture with Piz Daint and Tier-0 R&D extension

The nodes are then re-configured on the fly in order to satisfy all the requirements described in the previous sections. The site diagram with the Tier-0 R&D extension can be seen in Figure 2.

This commissioning turned out being very laborious: much extra tuning of the systems was needed, to address memory shortage and I/O load. On a scale of 10% of a full run, as for the particular ATLAS data stream *physics\_BphysLS stream*, a good performance was achieved, improving on the distributed spill-over turnaround. For a full run (*physics\_Main*) reconstruction, only half of the CPUs on a given node could be used. This due to the higher pressure on memory and I/O. The DVS layer revealed some software bugs hampering performance. Some were corrected quickly by Cray, some others fell out of the time frame of the project. In addition, an attempt has been made to use the Cray *Data Warp* technology [12] to instantiate *swap* on the nodes to help alleviate the memory shortage, but results have been inconclusive. It is believed that on the current system setup, the inefficiency observed could be recovered, at least partly.

## 5. Tier-2 operational experience

The full migration of the CSCS Tier-2 facility to Piz Daint was finalised by April 2019, with a scale up to 12k cores and the full phase-out of the Phoenix cluster. However, initial result in terms of CPU delivery and efficiencies have not fully met expectations. In particular, the model of a fixed node reservation shared by ATLAS, CMS and LHCb has turned out to be a bit weak, with several concurrent reasons causing a shortfall in the exploitation of the allocation. Some of the recurring issues met, include:

- Job scheduling complexity on the Piz Daint slurm creates challenges for: fair core allocation

**Table 1.** Reliability and performance indicators for Piz Daint (8 months) and Phoenix (12 months)

	<b>WC successful vs all jobs</b>	<b>CPU/WC efficiency successful jobs</b>
Piz Daint	73%	78%
Phoenix	68%	81%

between ATLAS, CMS and LHCb according to the target shares, continuity of allocation for ATLAS and capability of filling the allocated nodes.

- I/O pressure stressing the DVS service.
- In-RAM CVMFS cache is limited in size, so it is subject to instabilities when one or more experiments switch to make a more intensive use of CVMFS. This was the case when ATLAS switched to singularity with containers sourced off CVMFS.

In terms of reliability and performance Piz Daint scores similarly to the Phoenix cluster. For this evaluation, we measure for the two systems the Wall-clock (WC) efficiency based on successful over all jobs and the CPU/Wall-clock efficiency for the successful jobs (Table 1). The figures reflect the convolution of system and infrastructure inefficiencies in the production environment.

## 6. Conclusions

The Swiss HEP community is running a full WLCG Tier-2 on the HPC Piz Daint at CSCS. The complex integration process has been pioneered and led by ATLAS over the course of about five years. Since April 2019, the former commodity Tier-2 facility (Phoenix) has been phased out and the full Tier-2 runs on Piz Daint only, with about 12k statically dedicated cores shared by ATLAS, CMS and LHCb. The technical complexity of the system and some resulting integration solutions have resulted in increased operational pressure compared to the traditional cluster solution adopted in the past, and some issues are still open. Reliability and efficiency metrics for Piz Daint are comparable with the values of the Phoenix cluster. Some level of opportunistic usage of the HPC would compensate for some of the shortfalls in the allocation of resources we have observed, compared to what is reserved and should be available to the experiments. Other solutions are being explored, like dynamic allocation of nodes to each experiment to fulfil a yearly allocation. We should also consider the non negligible risk that when machines like Piz Daint are phased out and replaced by a newer generation, some integration solutions might be partly or fully invalidated. The considerations made above give us a clear indication that operating a fully pledged Tier-2 exclusively on a HPC does not provide tangible benefits and is operationally heavier compared to a traditional commodity system. As such, it is not the most optimised solution in view of the HL-LHC challenge and Swiss ATLAS is considering alternative schemes.

## Acknowledgments

We acknowledge the support of the Swiss National Science Foundation and thank CSCS for the provision of the integration systems.

## Copyright

2020 CERN for the benefit of the ATLAS Collaboration. CC-BY-4.0 license.

## References

- [1] Schmidt B *The High-Luminosity upgrade of the LHC: Physics and Technology Challenges for the Accelerator and the Experiments* Journal of Physics: Conference Series 898, **706** 2016 022002
- [2] Bird I, Campana S *WLCG Strategy towards HL-LHC*, 2018 CERN-LHCC-2018-019 ; LHCC-G-169
- [3] <https://www.cscs.ch/computers/piz-daint/>
- [4] Nilsson P et al., *The PanDA System in the ATLAS Experiment*, 2008 PoS(ACAT08)027.
- [5] Hostettler M *Enabling the ATLAS Experiment at the LHC for High Performance Computing*, Masterarbeit an der philosophisch-naturwissenschaftlichen Fakultät der Universität Bern 2015, CERN-THESIS-2015-391
- [6] Agostinelli S et al., *Geant4, a simulation toolkit* Nucl. Instrum. Meth. A **506** 2003 250-303
- [7] Blomer J and Fuhrmann T *A Fully Decentralized File System Cache for the CernVM-FS* Computer Communications and Networks (ICCCN) 2010 1-6
- [8] Sciacca F G, Haug S [ATLAS Collaboration] *ATLAS and LHC computing on CRAY* Journal of Physics: Conference Series 898, **082004** 2017
- [9] Sciacca F G, Weber M [ATLAS Collaboration] *Production experience and performance for ATLAS data processing on a Cray XC-50 at CSCS* EPJ Web of Conferences 214, **03023** 2019
- [10] Sugiyama S and Wallace D *Cray DVS: Data Virtualization Service* in Cray User Group Conference (CUG) 2008
- [11] Blomer J et al., *New directions in the CernVM file system* J. Phys.: Conf. Ser. 2017 **898** 062031
- [12] Liu N et al., *On the role of burst buffers in leadership-class storage systems* in Mass Storage Systems and Technologies (MSST), 2012 IEEE 28th Symposium on. IEEE 2012 1-11