




CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Commissioning CERN Tier-0 reconstruction workloads on Piz Daint at CSCS

HEPiX Fall 2018 – Barcelona, 10 October 2018

Gianfranco Sciacca, University of Bern (Speaker)

Pablo Fernandez, Miguel Gila, CSCS

10 October, 2018

Table of Contents

1. Motivation and background
2. Challenges
3. Tier-2 facility integration
4. Tier-0 spillover integration (ATLAS, CMS)
5. Results
6. Lessons learned

Motivation

- CSCS HPC resources are integrated with WLCG as part of the CSCS-LCG2 site
- CERN and CSCS have planned R&D projects for a deeper integration of the CSCS Computing resources with the LHC computing environment
(*in view of the challenges posed by the foreseen scale of the HL-LCH computing needs*)
- **”Development and testing of an infrastructure for accessing compute and storage resources in an HPC Centre”**
- **First use case proposed by ATLAS and CMS to CSCS:**
 - *Implementation of an environment supporting Tier-0 spill-over to Piz Daint*
 - **Goals**
 - Elastic provisioning of Tier-0 prompt reconstruction of the experiment RAW data
 - Support steady and on-demand spill-over / support computational peaks
 - Evaluate solutions and interaction in preparation for Run 3 (in 2021+)

Background

- **The Swiss HEP computing community and CSCS have started working on the HPC integration with the LHC experiment Tier-2 facilities in 2014**
 - **ATLAS Geant4 simulation**
 - Ran in production for 6 months on a Cray XK7
 - Integrated by means of a modified ARC CE, submitting remotely from Bern to CSCS
 - **LHConCray project (ATLAS, CMS, LHCb)**
 - Ran for about 2 years in 2016-17
 - Aimed at integrating Piz Daint with the LHC experiment frameworks
 - Targeted all experiment workflows (including user analysis)
 - Went in production with 1.6k cores in 2017
 - **WLCG Tier-2 facilities migrated to Piz Daint**
 - Decision taken at the end of 2017
 - >4k cores by April 2018, >10k by April 2019

Challenges

- **HPC is awesome**

- Powerful machines featuring top of the range technologies
- Piz Daint at CSCS: <https://www.cscs.ch/computers/piz-daint/>
- Benefits from economy of scales both in procurement and operation costs

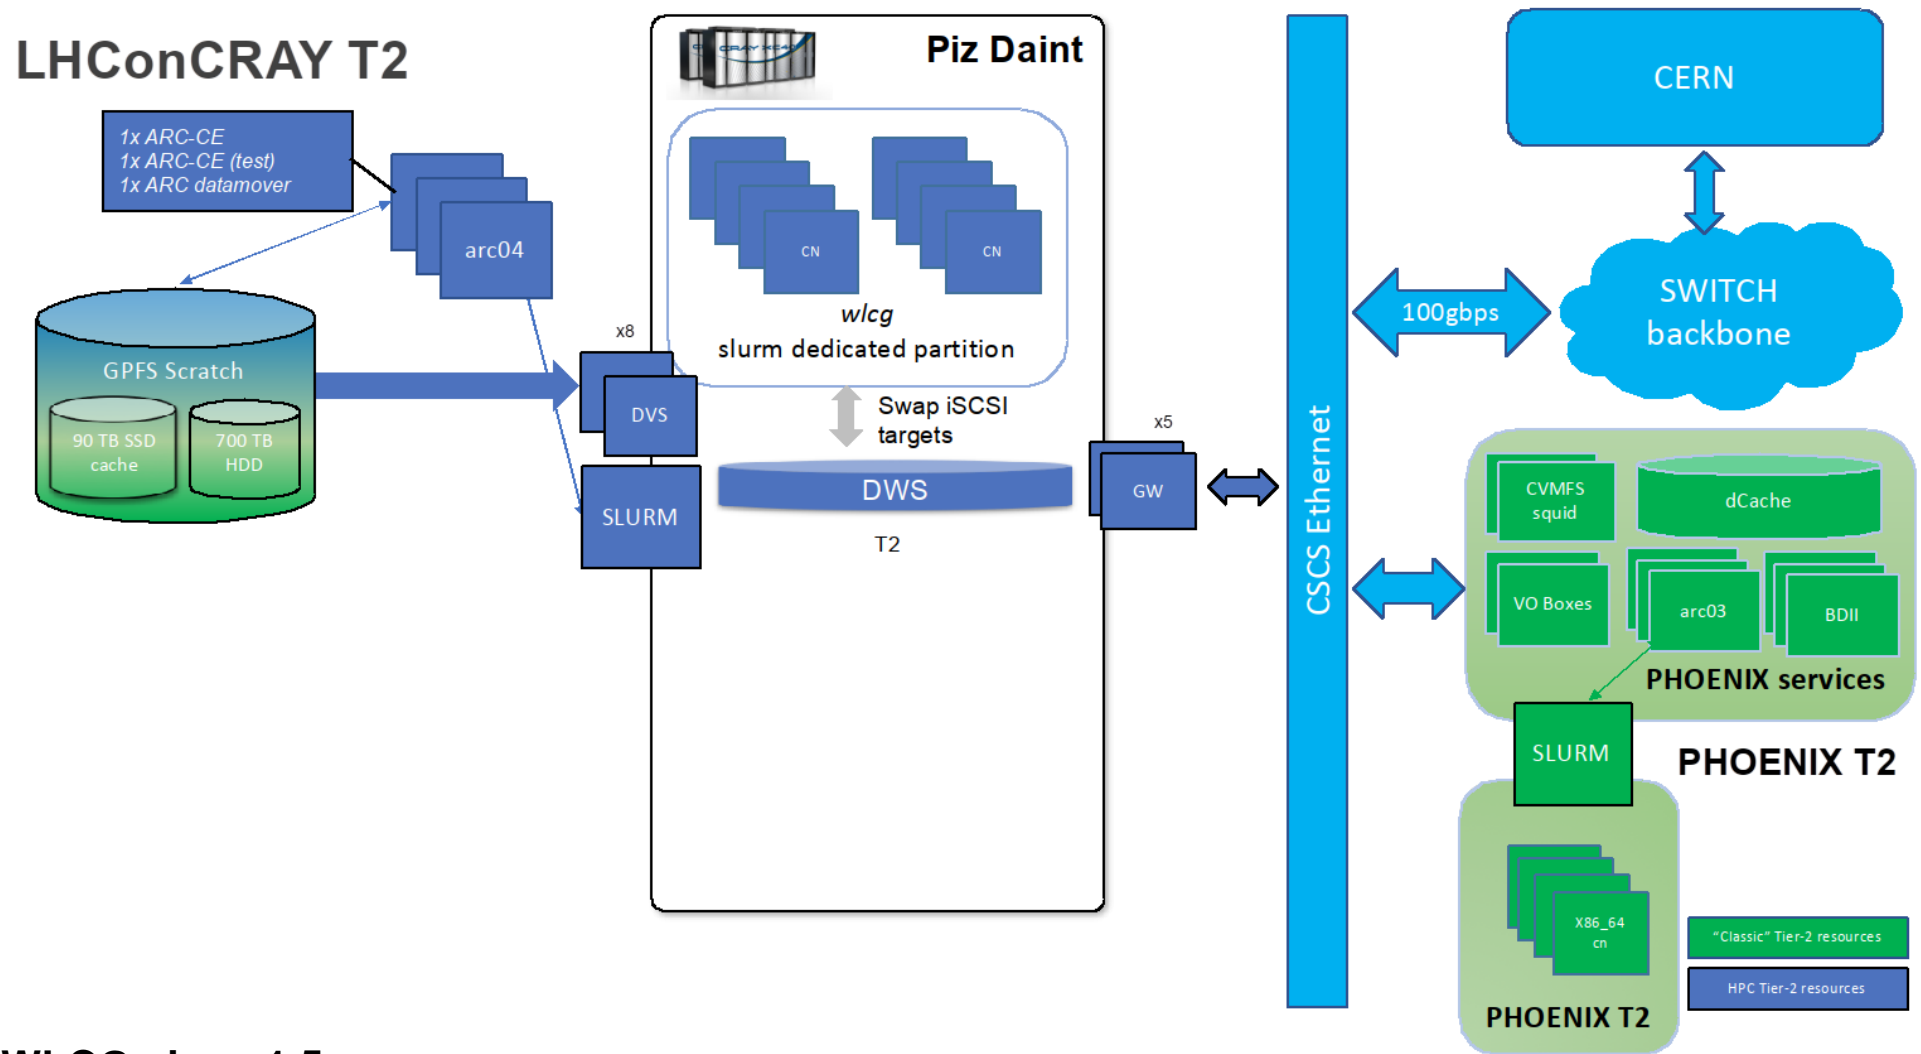
- **HPC is awkward** => *for HEP*

- Technologies and OS are optimized to accelerate parallel software
- Many expected linux tools missing
- Diskless nodes with 1 GB/core, some with 2 GB/core, no swap
- Network usage and I/O patterns not typical compared to standard HPC workflows

- **Container technology has gone a long way to make them look more like traditional linux systems**

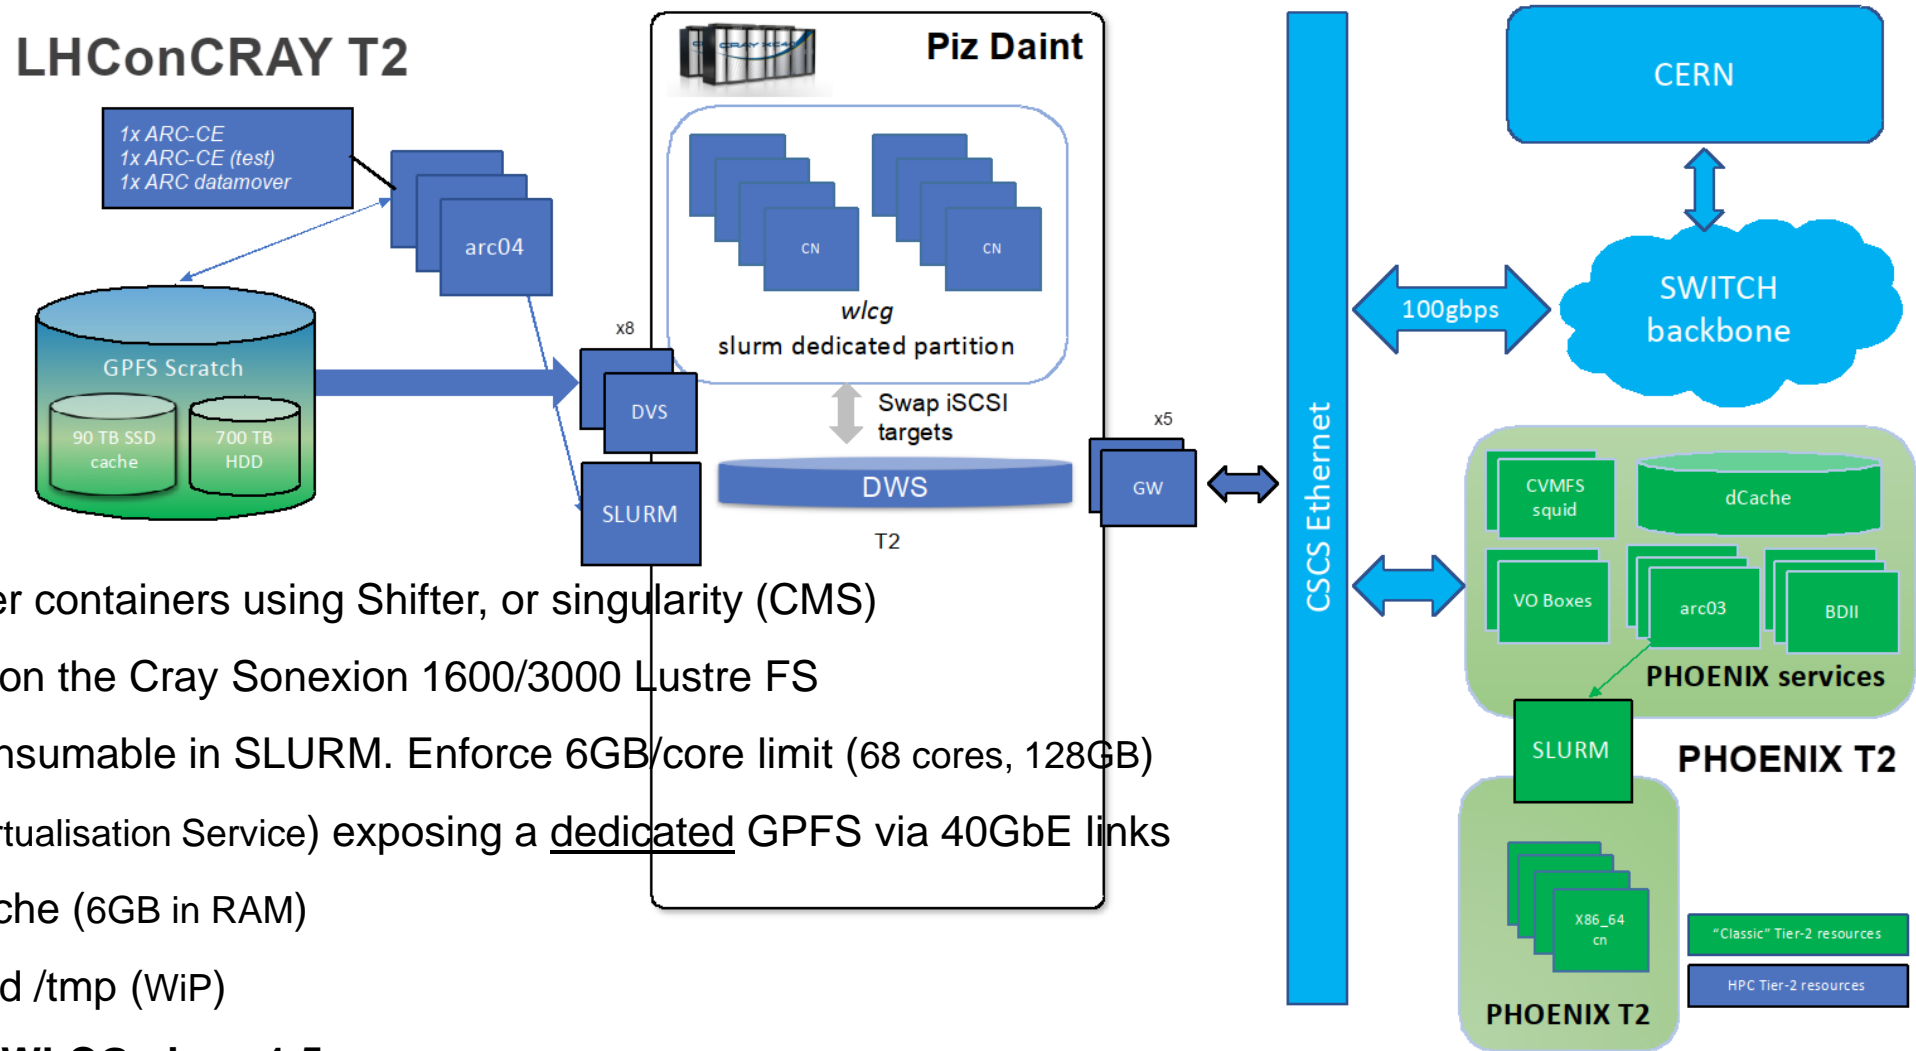
- But could not solve all challenges related to the integration with the complex LHC experiment frameworks

Tier-2 facility integration



- In production in WLCG since 1.5 years

Tier-2 facility integration



- Jobs run in Docker containers using Shifter, or singularity (CMS)
- Images (full WNs) on the Cray Sonexion 1600/3000 Lustre FS
- Memory is not consumable in SLURM. Enforce 6GB/core limit (68 cores, 128GB)
- DVS (Cray Data Virtualisation Service) exposing a dedicated GPFS via 40GbE links
- CVMFS tiered cache (6GB in RAM)
- DWS for swap and /tmp (WiP)
- **In production in WLCG since 1.5 years**



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich

From Tier-2 to Tier-0 integration

Tier-0 spillover integration

- **Extension of proven production Tier-2 design**

- Dedicated ARC servers and data-stagers
- GPFS as shared file system and ARC cache, with a dedicated SSD cache layer in front of it

- **No reserved resources, no node draining**

- ARC submits to a dedicated SLURM partition with up to 150 nodes (~10k cores) that overlaps with standard HPC jobs
- Can run multi-core or single-core jobs (up to 128GB RAM/64 cores per node). *Swap available.*
- Jobs from ARC get higher priority, taking over nodes as they become free (no idle resources)
- Nodes re-configured on the fly, CVMFS mounted on demand, when needed
- Allows for both Tier-0 *burst* and *steady* spillover to be processed without operational differences (*in principle*)

- **WLCG environment isolation**

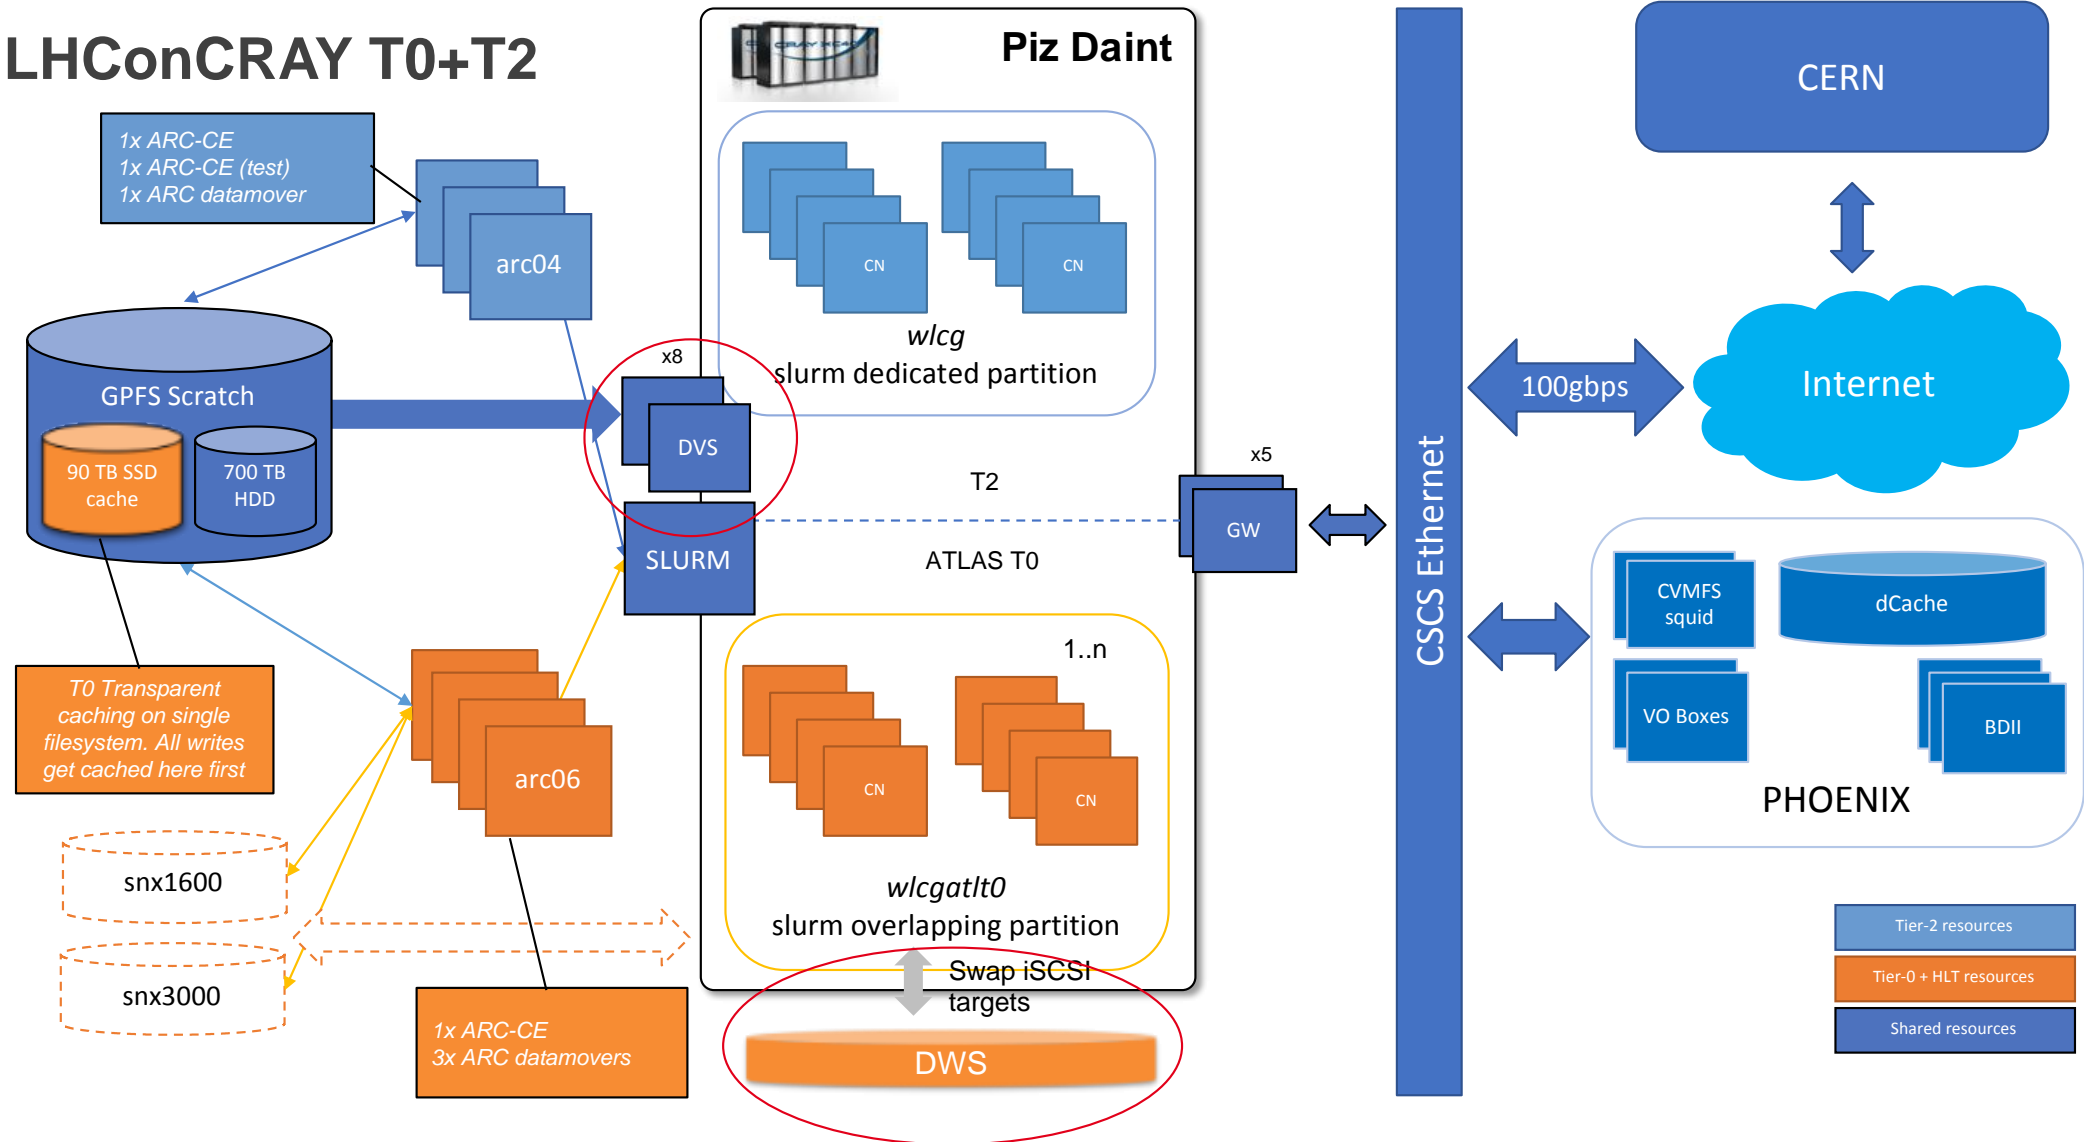
- Tier-2 jobs continue running (*on dedicated nodes*) along Tier-0 jobs

- **Excellent connectivity**

- 4x40gbps between Daint and dCache, 100 gbps link to Internet/CERN border

Tier-0 spillover integration

LHConCRAY T0+T2



Results

ATLAS workloads validation

- Dedicated PanDA queue, dedicated storage token
- Input RAW data copied asynchronously to the CSCS dCache, staged in via ARC
- Basics validation with HammerCloud tests (lightweight MC simulation jobs)
- **RAW data reconstruction on *physics_BphysLS stream*** $O(10\%)$ of *physics_Main*
 - Run continuously on $O(20)$ nodes (~ 1300 cores) - 0.7 TB input (710k events)
 - **Considerable tuning needed**: job corecount, OOM (solved by adding swap), cgroups and slurm config, CVMFS in-RAM, I/O (doubled the nr of DVS nodes), nr. of input events per job, PanDA brokerage, etc.

	Turnaround time	CPU/WC efficiency	CPU time/event	Wall time/event
Piz Daint	13h	32% (32-core)	20.8	78.4
Grid	18h (90%), 46h (100%)	53% (various)	13.6	33.5

ATLAS workloads validation

- **RAW data reconstruction on *physics_Main*** (input 35 TB, 32M events)
 - On demand runs on O(150) nodes (~10000 cores)
 - Memory needs slightly higher, swap needed, but failed to instantiate (Cray DWS bug)
 - Tried tuning corecount, could not pack the nodes
 - 3x16-core jobs performed better, but too many job transient files => DVS could not cope (another Cray bug)
 - Settled on 2x16-core jobs per node, run steadily, but large waste of resources



	Jobs per node	CPU/WC efficiency	Turnaround time	CPU time/event	Wall time/event
Piz Daint	2x32 core	27%			
Piz Daint	1x56 core	23% (effective 11%)			
Piz Daint	3x16 core	42% (effective 32%)			
Piz Daint	2x16 core	62% (effective 31%)	7d 7h	22.4	55.0
Grid	various	62%	5d 9h	18.2	31.6

ATLAS conclusions

- **We have (~)succeeded in validating Piz Daint for ATLAS Tier-0 spill-over**
 - Validated for *physics_BphysLS* stream
 - Higher WC/event due to high corecount used (serial portions of the workload)
 - Understood the current limitation for *Physics_Main* (Cray bugs)
 - Still not too far from the target
 - This has been a very laborious exercise that has involved several experts on both sides
 - Not surprisingly: we aimed at fitting a workload to a system that *does not entirely fulfil the hardware requirements*, so a lot of tuning needed be put in place
- **We consider validating Piz Daint for Tier-0 workloads an outstanding achievement**


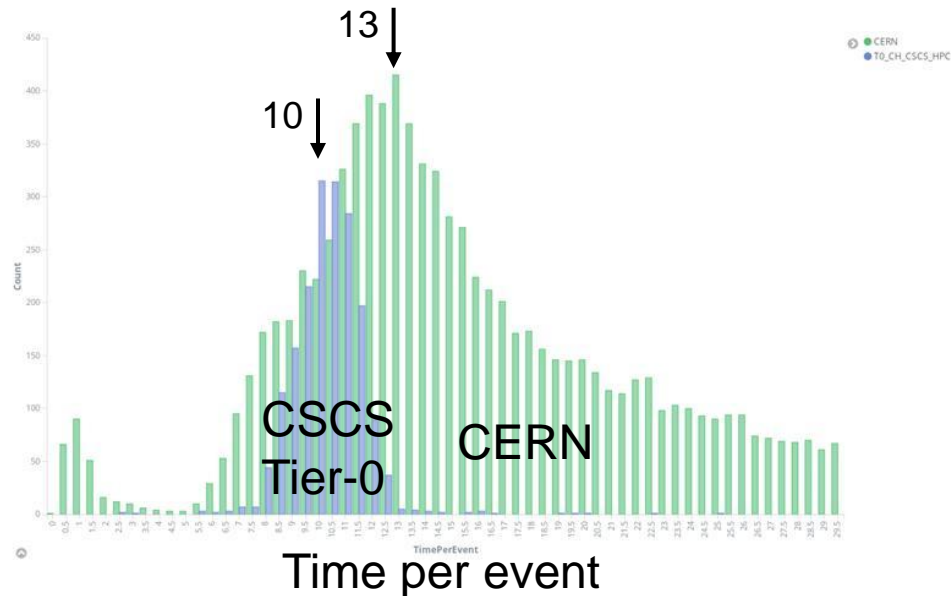
CMS health check

- Hammercloud commissioning – first check things are right
 - Initial issues with CVMFS stability and low efficiency, solved with cgroups configuration
- Production workflows – check the “site” can do regular Tier-2 prod
 - Few hundred cores, very low error rates, good time/event, good efficiency



CMS Tier-0 replay test

- ~2000 cores allocated at CSCS
 - Reconstruction jobs: processing 20% of a typical run with 8-thread jobs
- All data hosted at EOS @ CERN (local CSCS grid SE not used)
 - 13 TB input data read, 17 TB output data produced
 - Input read directly from CERN through the network
 - Output written to local scratch, then shipped to EOS at the end of the job



CMS conclusions

- Data to be taken with care, since the jobs running on both sites were not the same (still Tier-0 Reco)
- CPU efficiency is very much comparable for both sites
 - Slightly lower for CSCS, which can be explained by remote data read/write
- Time per event is difficult to compare
 - Still no evidence of significant performance difference
- We are probably ready to go into production.
 - We would need a longer sustained test, but we have no real evidence of huge problems

Lessons Learned


- We have shown that LHC experiments can use a general purpose HPC system transparently for all their workflows
 - Integration efforts were costly, first time Tier-0 workloads go to an HPC system
 - No major technical showstopper
- This required the centre to relax some policies
 - But nothing major against design decision or PRACE level policy
- We hope some HEP requirements will drive design of the next generation machine(s)



CSCS

Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

ETH zürich



Thank you for your attention.