Report on "CHIPP LHC Computing"

- Status of operation
- News on strategic and operational issues

Christoph Grab
ETH Zürich

28.08.2018

Status of operation

Overview Swiss LHC Computing Resources

- **Switzerland operates Tier-2 Regional Centres at CSCS and AEC (UBe)**
  - Switzerland is committed as member of WLCG to contribute resources; via MoU
  - resources provided to WLCG are exploited centrally by experiments
  - **Tier-2 operated at CSCS serves all 3 experiments: ATLAS,CMS, LHCb**
    ⟹ A) own linux cluster PHOENIX and B) shared resources on HPC/PizDaint
  - **Tier-2 operated at AEC-UNIBE serves ATLAS only**
    - Collaboration agreement for operation of T2 between CHIPP and CSCS/ ETHZ - with additional ETHZ funding (2007-2018)

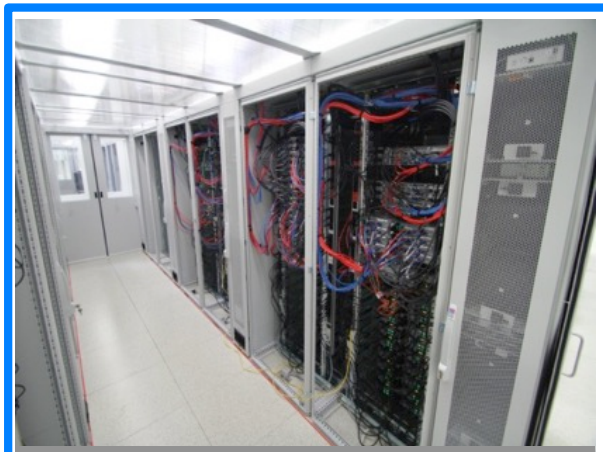

CSCS: Phoenix cluster

CSCS: Piz Daint-CRAY

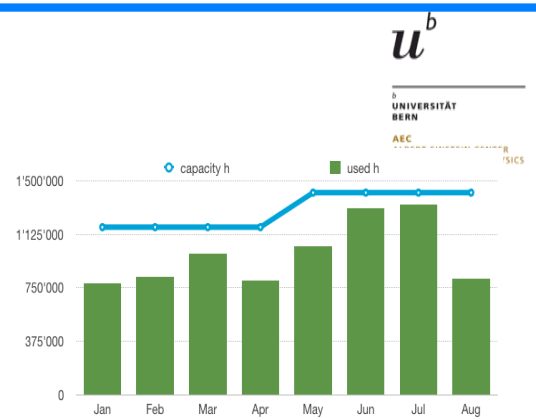

AEC: ATLAS T2

CSCS
Centro Svizzero di Calcolo Scientifico
Swiss National Supercomputing Centre

- **Complemented by local Tier-3 clusters** at PSI, UBe, UGe, UZH, EPFL

Overview Swiss *LHC* Computing Resources

EGI (European Grid Infrastructure); NGI

WLCG (Wordwide LHC compute Grid)

CSCS T2 -ATLAS, CMS, LHCb
- **PHOENIX:**
  - 5800 cores; 65 kHS06
- **PizDaint**:
  - 3876 cores; 46 kHS06

Common: ~4100 TB disk

CPU / Disk share~ 40:40:20 / 43:40:17

- 80 Gb/s to 100 Gb/s backbone

AEC-UNIBE T2/T3 (ATLAS)
- 2400 cores ; ~22 kHS06

1000 TB disk / 250 TB scratch

2x10 Gb/s to backbone

PSI-ETHZ-UZH T3 (CMS) 0818

1184 cores;13.4 kHS06;1168TB

+ 229 RO cache

DPNC-UNIGE T3 (ATLAS)

784 cores; 6.0 kHS06; 828 TB

+ add 110 cores on UGe Cluster

UZH T3 (LHCb) 0818

560 cores; 6.8 kHS06; 340 TB

incl. 88 cores on ScienceCloud

EPFL T3 (LHCb)

410 cores; 6.8 kHS06; 80 TB

Networking by SWITCH)

Coordination

CHIPP Computing Board / +NGI/EGI

Regular F2F meetings of CCB

Operation : Monthly meetings; (CH; EGI/GDB). CG_0818

Note: sum of Tier-3 resources [~33 kHS06; ~2.4 PB]
surpasses ~1/2 of Tier-2 resources (except ATLAS)

CHIPP resource statistics (7.2017-7.2018)

PHOENIX AND PizDaint CPU

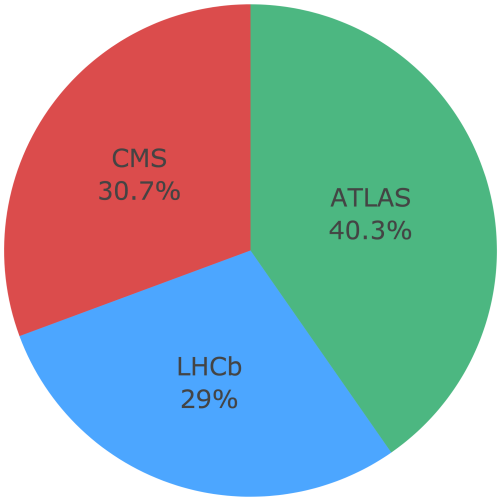

ATLAS

CMS

LHCb

dCache Storage available

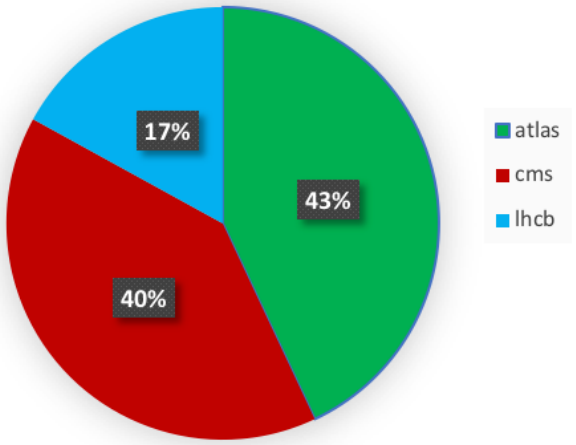

Storage is common for Phoenix + PizDaint

Total of 4.1 PB

Resources delivered to WLCG:

in Aug 2017 to Jul 2018
for WLCG a total of

CHIPP delivered to WLCG ~70 MHS06*hours

- CSCS : ~60 MHS06*hours (42 Phoenix; 17 HPC)
- ATLAS-Bern ~10 MHS06*hours

CHIPP delivered resources (2018)

Accounting in CPU-hours (PizDaint, Phoenix)

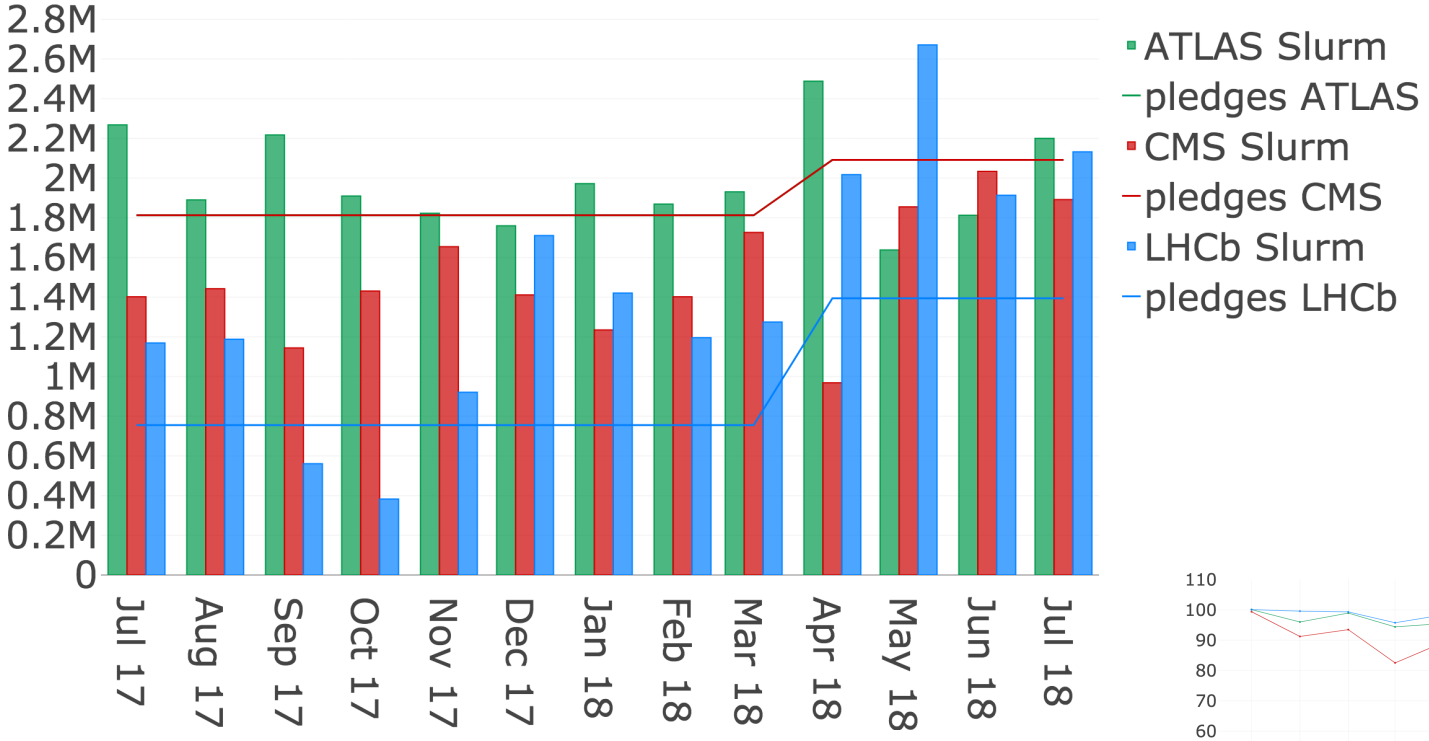

Integral over the year (8.17-7.18)
⟹ CHIPP/ T2@CSCS delivered to WLCG ~60 MHS06*h

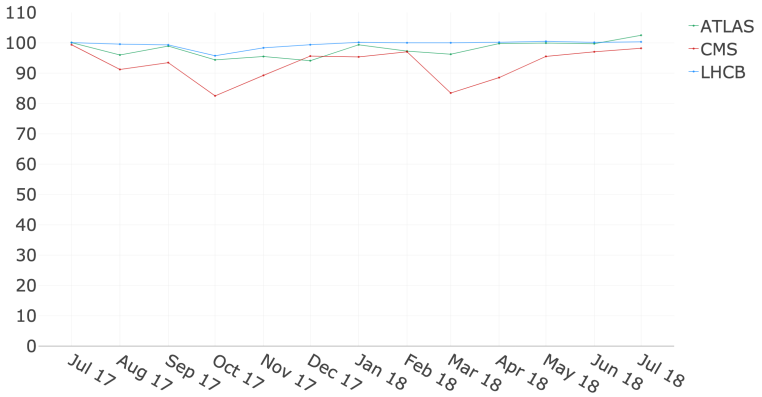

CPU efficiency EGI in %

CHIPP WLCG Computing Resources (8/18)

CHIPP Tier-2 pledged vs delivered resources (source REBUS)

Installed Capacities

Year: 2018 Month: 8

Infrastructure	Site Name	Physical CPU	Logical CPU	HEPSPEC06	Disk (GB)	Tape (GB)
EGI	CSCS-LCG2	442	10,084	116,064	5,162,057	0
EGI	UNIBE-LHEP	147	2,278	23,970	1,070,870	0
Total		589	12,362	140,034	6,232,927	0

Showing 1 to 2 of 2 entries

Status 8/2018

Federation Pledges

Year: 2018

Pledge Type	ALICE	% of Req.	ATLAS	% of Req.	CMS	% of Req.	LHCb	% of Req.	SUM	% of Req.
CPU (HEP-SPEC06)			54,000	5%	36,000	4%	24,000	17%	114,000	5%
Disk (Tbytes)			2,100	2%	1,600	2%	800	14%	4,500	3%

Showing 1 to 2 of 2 entries

https://wlcg-rebus.cern.ch/apps/topology/federation/259/

➔ Pledges to WLCG for 2018 are met; next pledges for 2019 are due in Sep. 2018.

Swiss Tier-3 resources (8/18)

Swiss Tier-3 resources are undispensible tools
and exist in quite different “flavours” for :

- ATLAS: each at UBern and at UGe
- CMS: common T3 for ETHZ, UZH, PSI at PSI
- LHCb: each at UZH and EPFL.

- Their capacity sum up to ~50% and 70% of CPU and storage of the Tier-2 (at CSCS w/out AEC) .

Swiss National Network

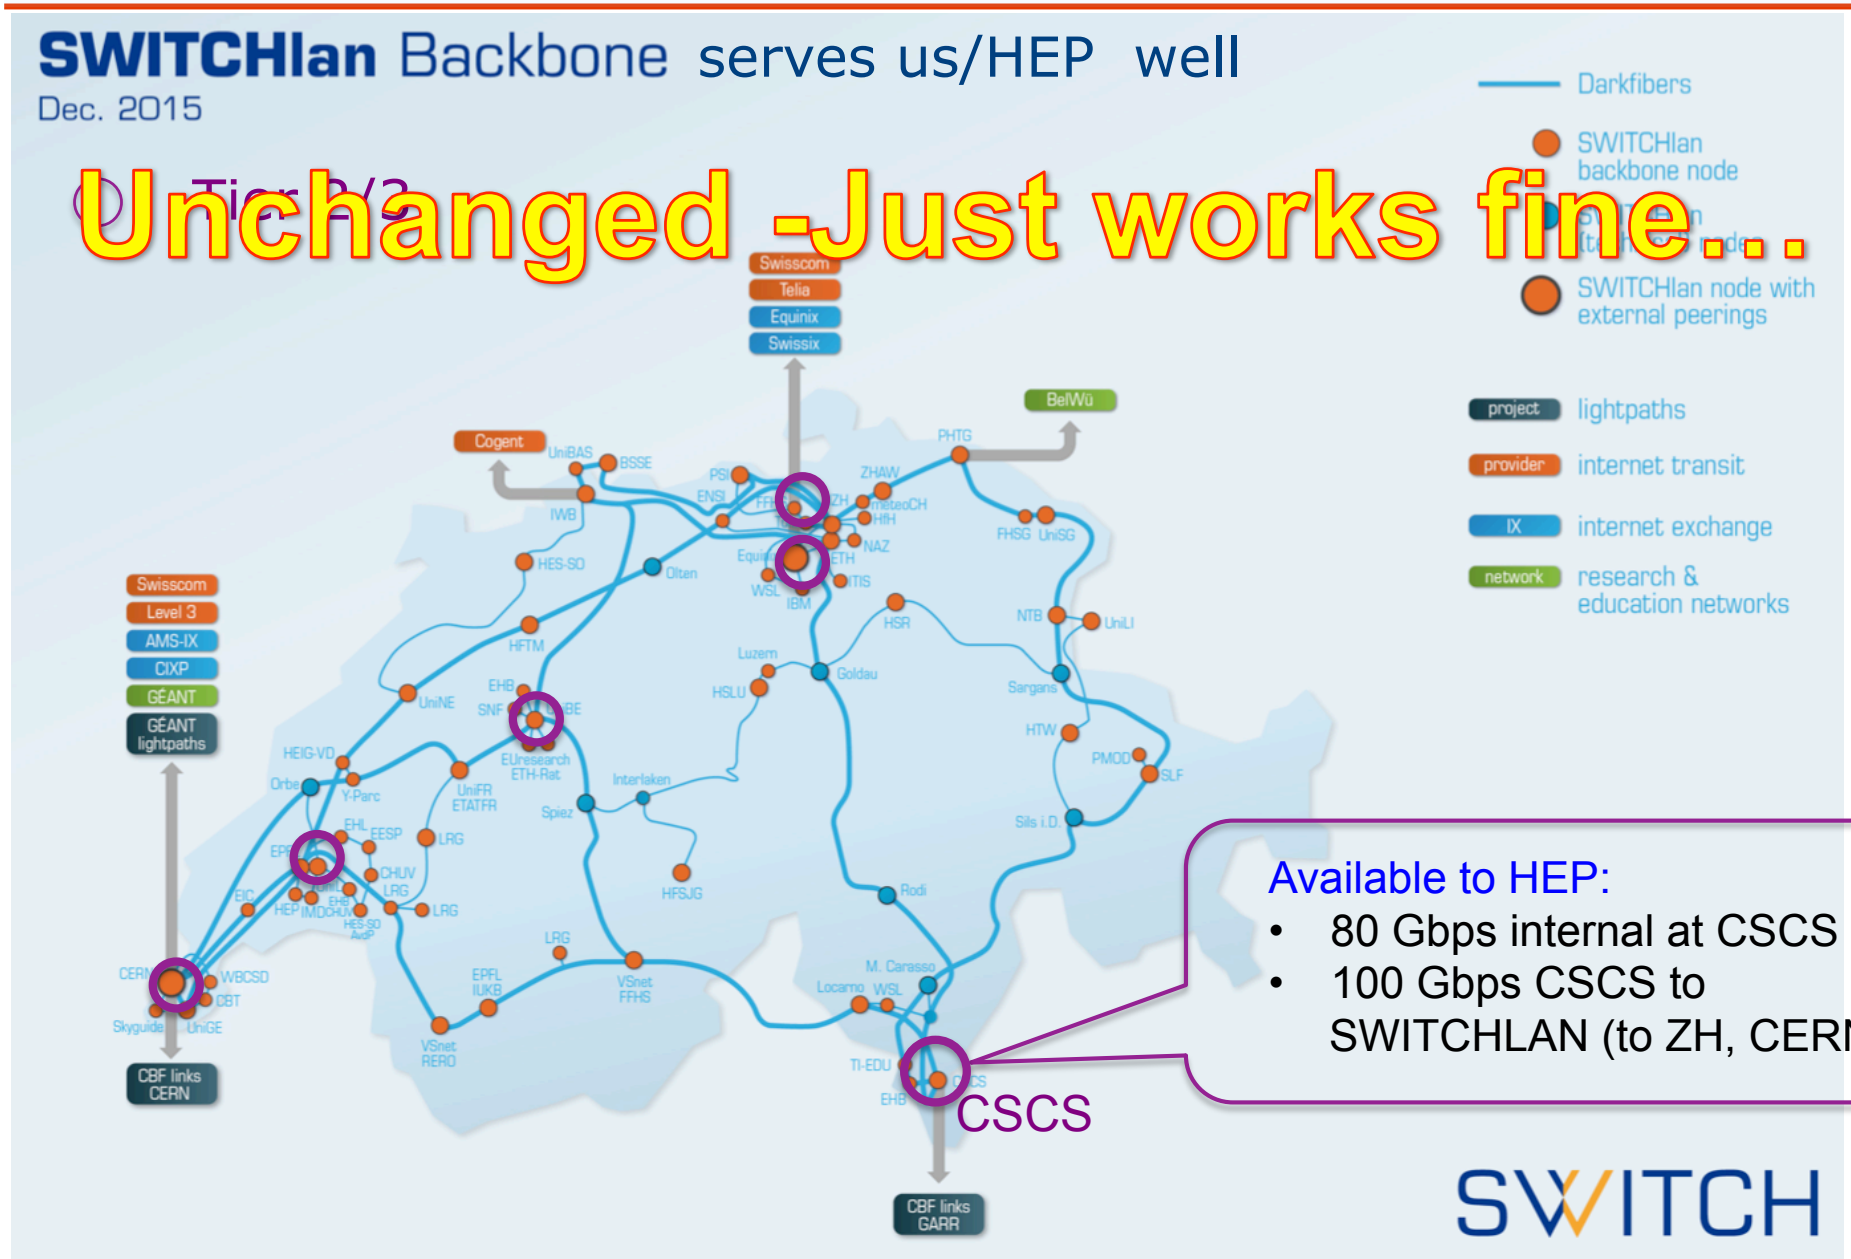

News on
Strategic and Operational Issues

I. Collaboration Agreement ETH – CHIPP

- CHIPP – CSCS cooperation(to provide resources and deliver our pledged resources to WLCG) is fixed in a "Collaboration Agreement" between ETHZ and CHIPP.

- This CA describes the duties and rights of both parties:
  - The CA was originally negotiated 2007 and renewed in July 2013;
  - **The CA has now (July 2018) been successfully renewed for another four years, for the duration 1.1.2019-31.12.2022 !.**

- In the CA, ETHZ also commits additional voluntary funding by ETHZ/CSCS, which covers costs for electricity & cooling and additional personnel (1 FTE)

II. Challenging and strategic issues (1)

- Challenge for LHC computing for the HL-LHC era in the 2025+ :
  - Present LHC compute models do not scale beyond ~2025,
    computing needs are expected to go up ~50 x in around 8-10 years;
    technology advance alone will accommodate a factor ~5
    → **need roughly another order of magnitude increase**

- Many activities exist worldwide to face this HL-LHC challenge

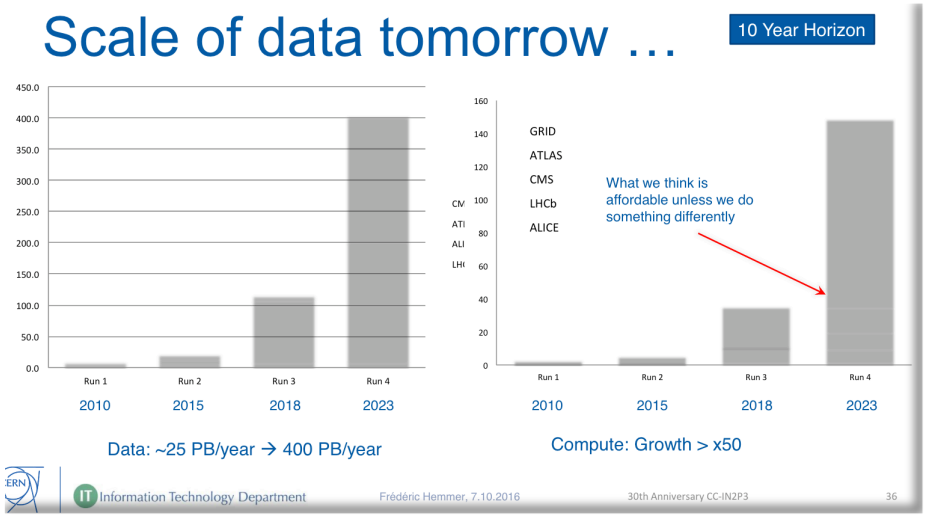

- Adapt to new architectures (HPC, ARM, … mobile??)
- exploit all existing resources
- BUT: Invest also in SOFTWARE,
  i.e. algorithms … (see "Community White Paper" by HEP-SF)
  → means also in PEOPLE !

II. Strategic + Operational issues (2) – LHC@HPC

- Switzerland started project LHConCRAY in 2016 (initiated at AEC-Bern) to test possibility and economy of LHC workloads on HPCs.

- ***In December 2017: concluded tests successfully.***
  - ***Team CSCS+CHIPP succeeded to run ALL LHC job-types on CRAY ! found similar job efficiency as PHOENIX, but higher economic value***
  - ***Meeting of "CHIPP LHC computing board" on 7.12.2017, decided to transition to HPC for providing the Swiss T2-resources @CSCS.***
  - **This means:**
    1) CHIPP will continue to provide the pledges of Switzerland towards WLCG
    2) CSCS will provide shared HPC resources for LHC computing, based on same FLAT budget by FLARE/SNF (and ETHZ+Uni contributions)
    3) PHOENIX as a "separate dedicated cluster" will be phased out.
    4) AEC at Bern continues providing additional ATLAS-T2 resources

II. Strategic + Operational issues (3) Resource Strategy of CHIPP Tier-2 @CSCS

- The growth growth of resources of Tier-2 @CSCS will be provided using shared resources at CSCS
  (e.g. yearly increase expected: compute ~27% and storage ~15%)

- Transition from "HW investment model" to "Paying for services" model : **"LCaaS" (LHC-Computing as a Service).**
  CHIPP will pay for resource provisioning and operational costs:

  - Compute : Dedicated number of nodes, allocated to CHIPP
  - Storage : Allocation of disk space via the SAN (as today)
  - Operational costs: ie. personnel, operations, network, licenses, …

III. Strategic Issue: Pilot development projects (I)

A) Pilot projects testing Tier-0 spill-over of data to CSCS:

exploit fast network and flexibility of HPC@CSCS resource allocation.

- **ATLAS/CERN: (Sciacca, Klimentov etal):**
  Test of online data-transfer and reconstruction of events. It was challenging and laborious job. However was finally successfully scaled up to 150 nodes (10'000 cores).

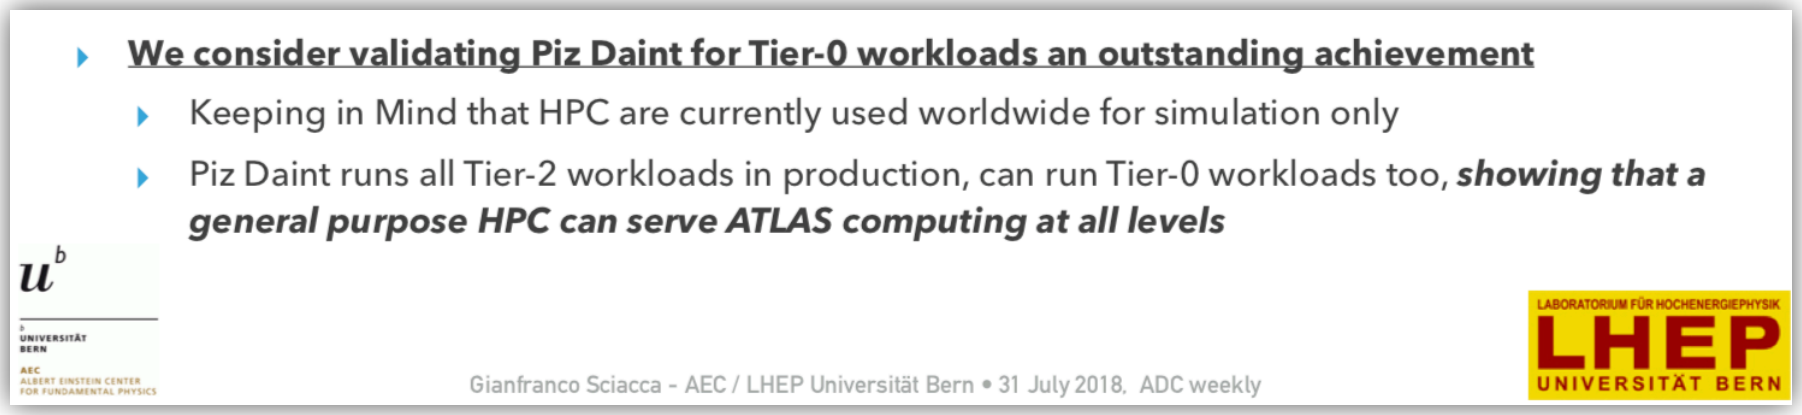

- **CMS/CERN: (Boccali, Wissing, et al):**
  Tests to run standard CMS Tier-0 workflows directly at CSCS were very successful, including stage-in-and-out to CERN; scaling to 10k cores still missing
  Next: Test scaling.

III. Strategic Issue: Pilot development projects (II)

Process of setting up cooperation CSCS – CERN to "Use of Remote Swiss HPC Facilities"
in view of solving the computing challenges of the future for HL-LHC

B) Pilot projects using GPU/CPU (in progress):

Goal is to exploit the GPU resources on HPC@CSCS through direct remote access (partly transparent).

Various authors from the ML community at CMS/ETH, with different physics applications (Pierini, Musella, Kasieczka, Pata, Vlimant, etc)

- CMS/ETH : Improving Precision Through Deep Learning in HEP
- CMS/CERN Event Generators based on Deep Learning for High-Luminosity LHC
- CMS/CERN: Developing Deep Learning Applications for the CMS HGCAL-Calorimeter
- CMS/CERN: Applications of Deep Learning in High Energy Physics

⟹ Some projects successfully completed (and used in publications), some still in progress. ⟹ **extremely useful for certain HEP applications.**

https://twiki.cern.ch/twiki/bin/viewauth/CMS/CmsMLProposalsCSCS

Organisation for CHIPP computing (III) (08/2018)

CHIPP LHC Computing board decides on strategic issues. It is advised by technical a advisory board.

Technical Advisory board

G.Sciacca (UBe)
T.Golling (UGe)

C.Grab (ETHZ) chair
T. Kljinsmaa (ETHZ)
D.Feichtinger (PSI)
N.Loktionova (PSI)

R.Bernet (UZH)
L.Pescatore, (EPFL)

P.Fernandez, M.Gila,
M. De Lorenzi (CSCS)

CHIPP LHC computing board

M.Weber (UBe)
T.Golling (UGe)

C.Grab (chair, ETHZ)
G.Dissertori (ETHZ)
D.Feichtinger (PSI)
F.Canelli (UZH)

N.Serra (UZH)
A.Bay (EPFL)

T. Schulthess (CSCS)