Report from the CHIPP Computing Board on Swiss LHC Computing

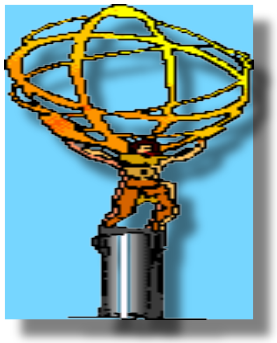

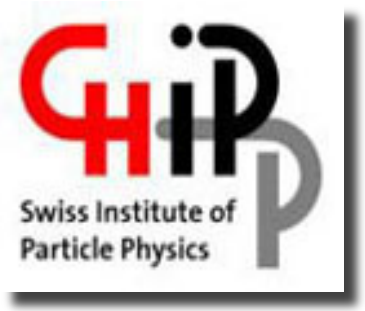

Christoph Grab

CHIPP plenary, July 1, 2015

Status of Operation of Swiss LHC Computing within WLCG

"LHC Computing" in CH

- **Switzerland operates a Tier-2 Regional Centre at CSCS**
  - Compute-cluster integrated into "Worldwide LHC Computing Grid" WLCG.
  - Switzerland is committed as full member; signed MoU; contributes resources.

- **Serving all three experiments :** ATLAS, CMS, LHCb
  - Collaboration agreement for operation of T2 between CHIPP and CSCS/ETHZ (2007-2012; 2013-2018 renewed and ETHZ funding secured)
  - Presently: ~37 kHS06; ~2100 TByte disk; (~2–4% of whole WLCG)
  - Tier-2 supplemented by dedicated ATLAS federation resource at AEC-UNIBE

Swiss Tier-2 Phoenix cluster at Lugano

- **Complemented by local tier-3 clusters** at PSI, UBe+UGe, UZH+EFL

Swiss LHC Computing Infrastructure

EGI (European Grid Infrastructure), **WLCG** (Wordwide LHC compute Grid)

EPFL T3 (LHCb)
480 cores; 3.6 kHS06;140 TB

UZH T3 (LHCb)
200 cores; 3.2 kHS06;200 TB

PSI-ETHZ-UZH T3 (CMS)
626 cores; 7.7kHS06; 760 TB
•https://wiki.chipp.ch/twiki/bin/view/CmsTier3/WebHome

DPNC-UNIGE T3 (ATLAS)
784 cores;4.3 kHS06; 530 TB
Direct 10 Gb/s to CERN IT

CSCS T2 (ATLAS, CMS, LHCb)
- 3700 cores; 37 kHS06
- ~2100 TB disk
- 10 Gb/s to 20 Gb/s backbone

•https://wiki.chipp.ch/twiki/bin/view/LCGTier2/WebHome

AEC-UNIBE T2 (ATLAS)
- 2500 cores; 15 kHS06;
- 450 TB disk
- 10 Gb/s

Operation
- Monthly meetings; (CH; EGI/GDB).

Networking by SWITCH)

Coordination
CHIPP Computing Board
Regular F2F meetings of CCB
(also SwiNG as NGI interface to EGI)

CG 6/15

Cluster Performance (2012-5.2015)

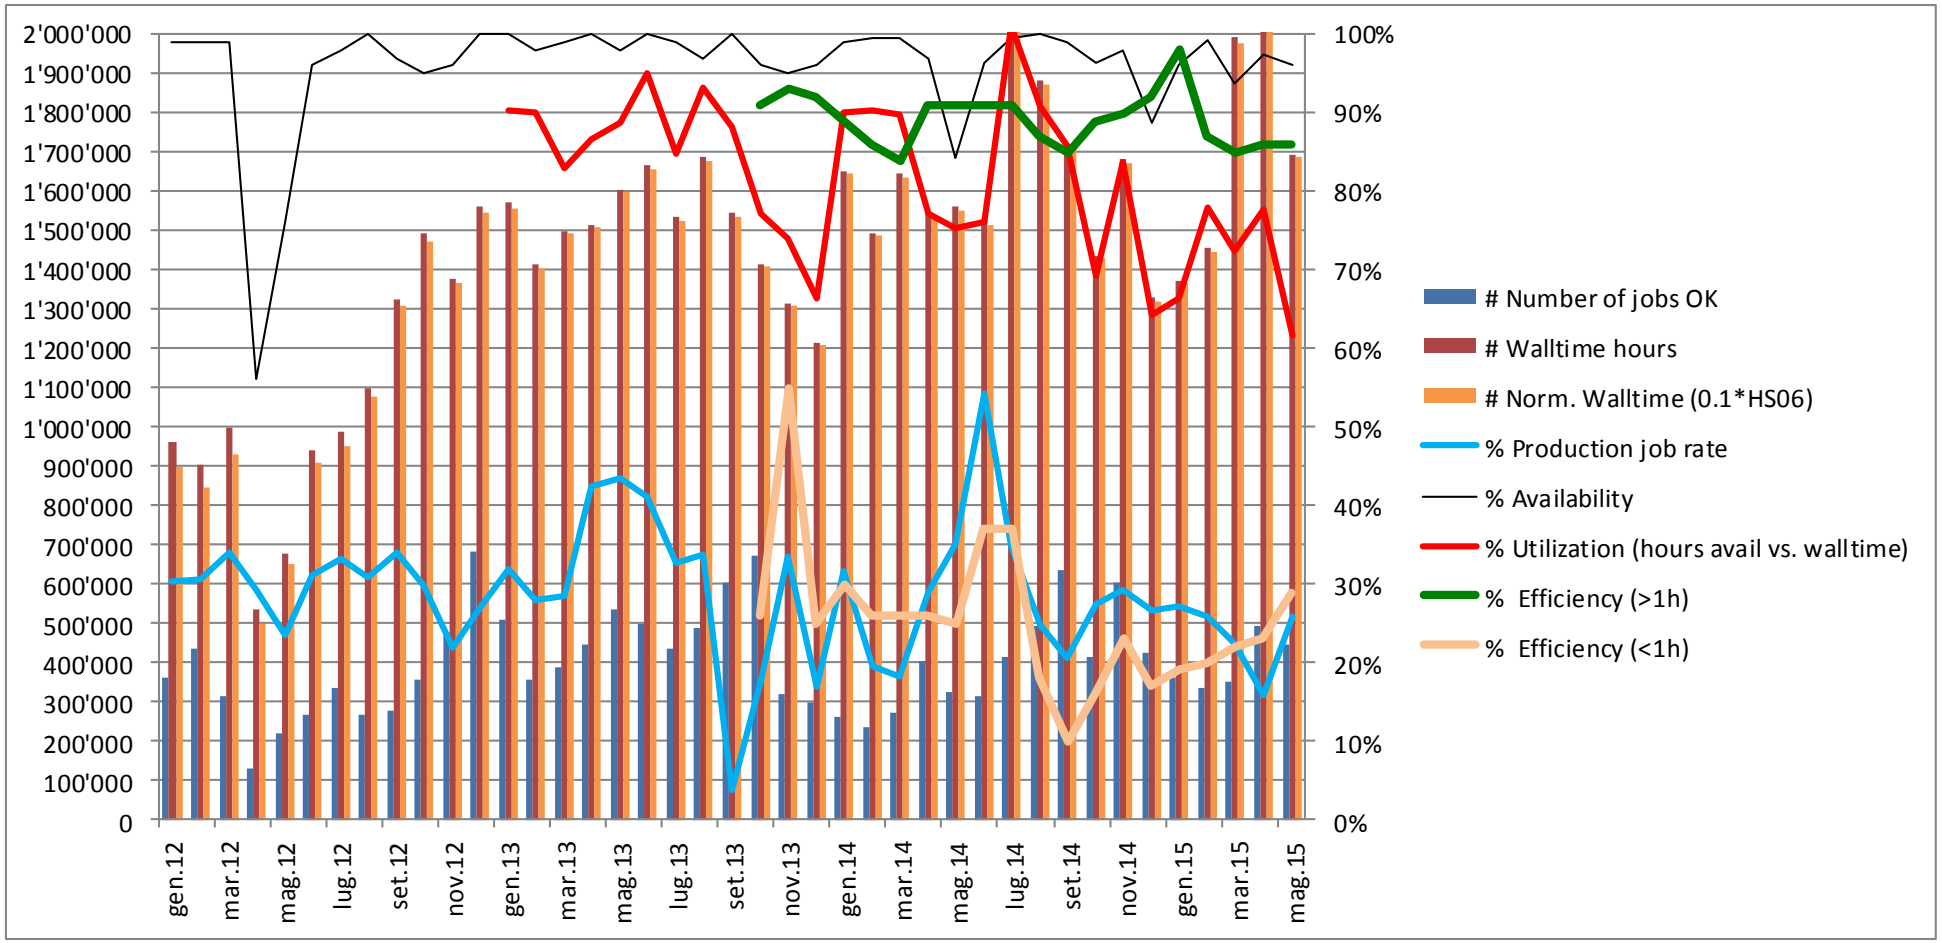

Overall high availability (>95%) and efficiency typ. 85-95% achieved !

(dip in Apr/May 2012 due to move to Lugano)

- https://wiki.chipp.ch/twiki/bin/view/LCGTier2/WebHome

Usage: Swiss T2@CSCS in 06.2014–05.2015

Swiss Tier2 CPU shares

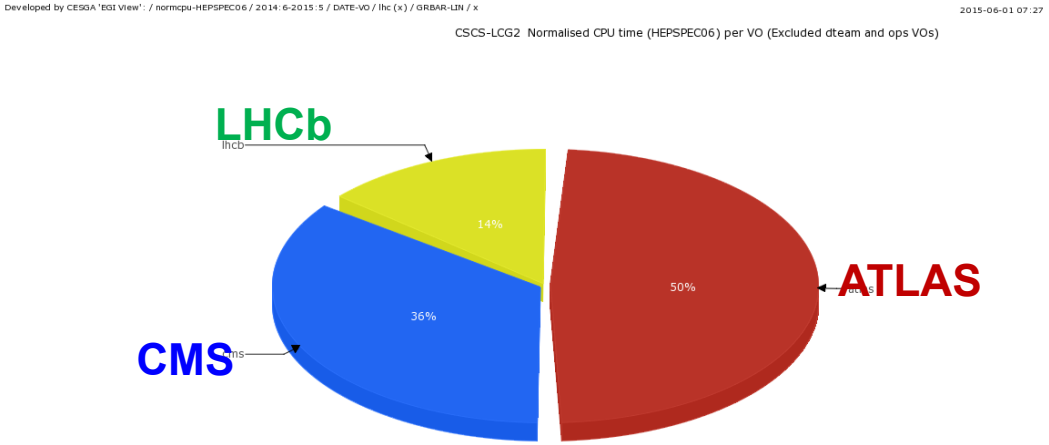

Swiss Tier2 Disk shares

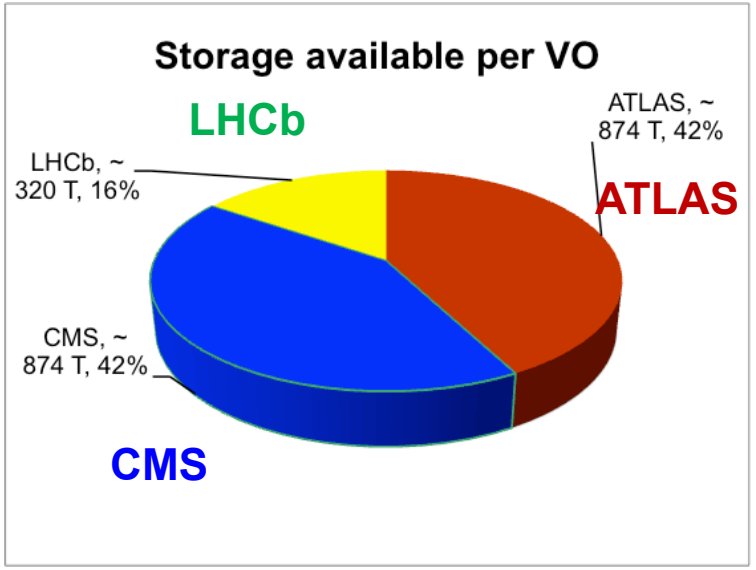

Worldwide Tier2 CPU usage

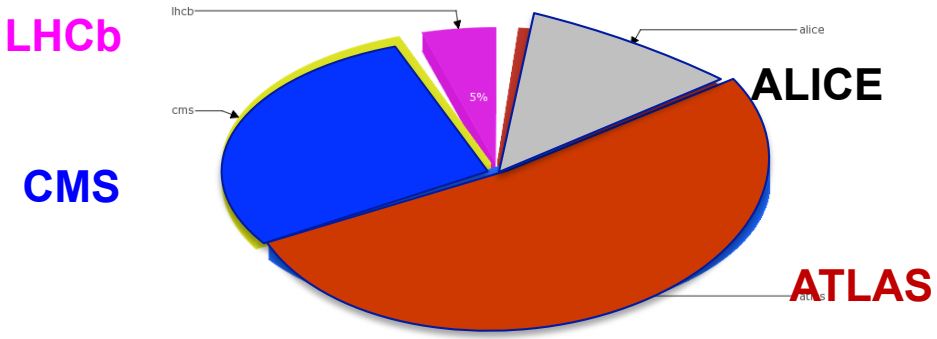

http://accounting.egi.eu

- CSCS fairshare ratio 40:40:20
- effective CPU usage: 40:32:28
- CSCS disk ratio: 42:42:16

Swiss Tier-2 Resources Evolution

Phoenix Evolution for 2007 – 2017 in CPU, Storage and Power

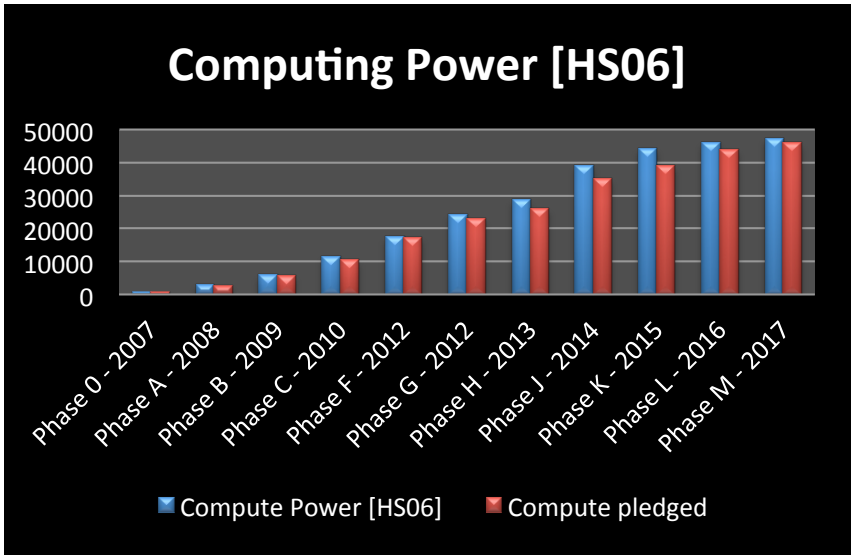

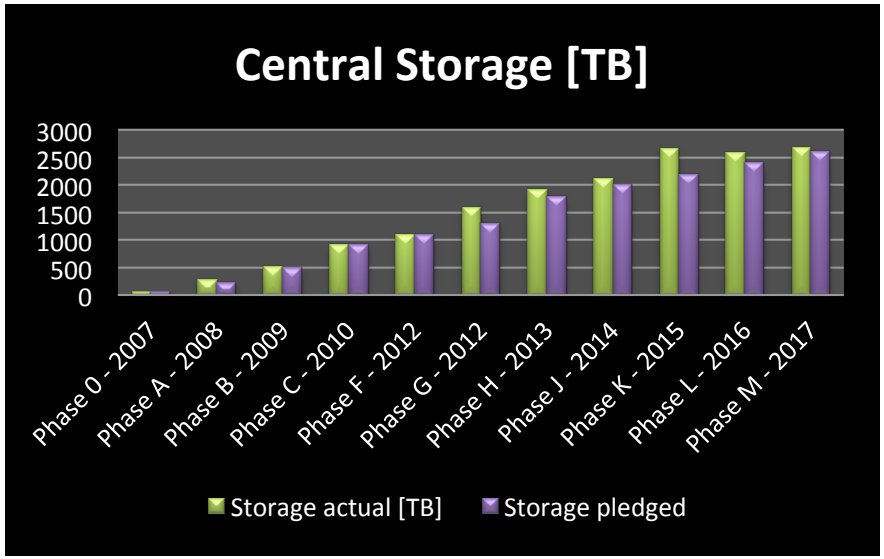

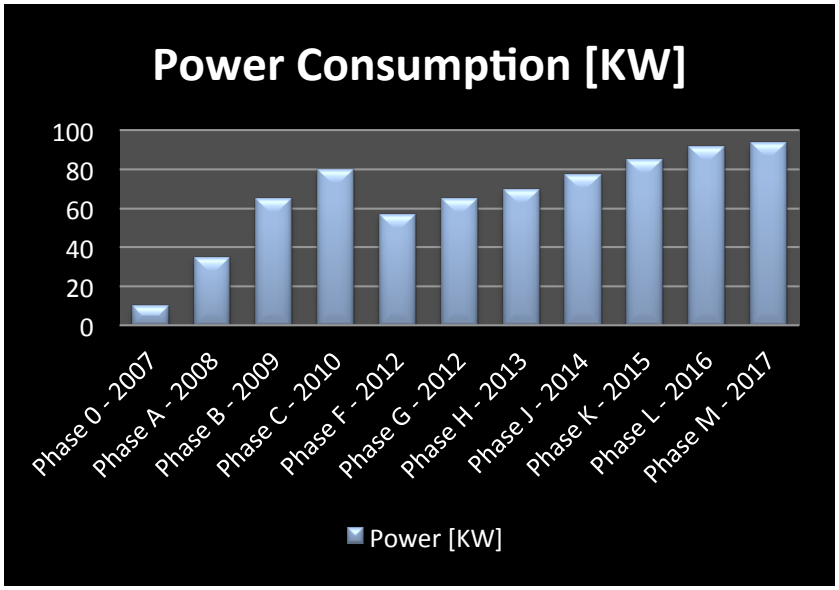

Swiss Tier-2 Resource Planning

Summary of resources provided and planned for Swiss Tier-2 at CSCS.

Resource	2012 (delivered)	2013 (delivered) Phase G	2014 (delivered) Phase H	2015 (projected) Phase J	2016 (future) Phase K	2017 (future) Phase L
CPU (kHS06)	17.4	23	26	35	40	44
Effective disk (TB)	1090	1300	1800	2300	2600	2900

Updated - June 2015

Phase K: financed; to be implemented in 2015; meet pledges in 1.4.2016.

Phase L: planned for implementation in 2016 to meet pledges of 1.4.2017.

More details in backup.

Swiss T2@CSCS Resources Planning

Planning assumes: flat budget

- **Planned investments (FLARE/SNF) of T2 per year**:
  - **HW** replacements and additions based on RRB recommendations of "flat budget". Provides about 20% increase of resource "power" per year (due to technical improvements).
    - ➔ investment for hardware for ~650 kCHF / year

- **Personnel for operation at CSCS**:
  - 1.5 FTE to support Tier-2 operation at CSCS, covered by FLARE
    Note: salary cap by SNF for technicians will be alleviated through supplementary grants, courtesy SNF.
  - 1 additional FTE is covered by ETH internal funds

- **Other resource items**
  - Recurring power/infrastructure costs (~100kCHF) carried by ETH
  - Tier-3 hardware costs covered by institutes
  - Tier-3 manpower covered by inst+experiments (Uni. and MoU)

Adaptation to LHCb compute models

- LHCb had modified its computing models – operating some Tier-2 also as analysis centre
- Swiss T2/ we adapted our T2 allocations

- Presently LHCb was allocated ~320TB (16%) and their “normal” CPU shares ~7 kHS06 (20%)

- LHCb now operates within the new computing model.

Swiss ONLY resources at CSCS

- **There are "Swiss only resources" available at CSCS for exclusive use by Swiss users (not pledged for common experiment usage)**

- **Currently (Q2/2015) consisting of**
  - **375TB of NetApp storage (**2 controllers, 2 servers, total 120 disks)
  - **8x HP WNs (~3.5 kHS06, 40 cores/node, 128 GB/node, total 320 core)**

- **Present implementation** (may change based on needs and feasibility)**:**
  - **Storage** is added to Central Storage (dCache) and enabled to users.
  - **Compute** nodes are running jobs,
    Swiss users have increased priority based on special VOMS mappings.

  - **Storage and compute distribution** according to :
    - 40:40:20 for compute (configuration is identical to non-exclusive CH resources)
    - 44:44:12 for storage (ATLAS ~165TB, CMS ~165TB, LHCb ~45TB)

Tier-3 Resources in CH

Swiss Tier-3 resources exist in:

- ATLAS: each at UBern and at UGe
- CMS: common T3 for ETHZ, UZH, PSI at PSI
- LHCb: each at UZH and EPFL.
- They (w/out AEC) sum up to ~50% of T2@CSCS

- Crucial for efficient, fast and flexible analyses in their final stages, **Indispensable tool** to do competitive physics
  ➔ they need continued support by Unis and SNF.

- Close links and communication to our Tier-2.
- Numbers included in the summary tables
- Details given in backup

Other items

Comments on Grid Certificates

- Service for granting grid certificates done sofar by SWITCH will stop on 30.9.2015.
- Chosen solution: we will get certificates via CERN or EGI.
  - CERN users: continue to get the *user* certificates from CERN
  - For all Non-CERN users:
    - AEC at UniBern will act as Swiss Registration Authority (RA) towards the EGI Certificate Authority (CA)
    - All server/*host certificates and user* certificates can be requested via: http://www.lhep.unibe.ch/sits/certificates.html
    - Note: SwiNG is the Swiss EGI member, so in principle only SwiNG member institutions are eligible (special cases may also apply)

Further items to note …

- Evolution of EGI secured : EGI-Engage granted by H2020 ✔
  https://www.egi.eu/about/egi-engage/
- Swiss NGI link to EGI secured through “NeI-CH” project, approved for 2 years by CRUS (Swiss-universities) ✔
- Significant efforts ongoing by experiments to live within realistic computing budgets; e.g. ✔
  - Increase efficiency of software (simulation, reconstruction..)
  - Increase efficiency of existing resource usage (data placement ..)
  - Investigate exploitation of other architectures (HPC …) etc.
- “HEP software foundation” setup to attack common problems:
  - Common themes in terms of SW knowledge, licensing, teaching, technical fora…

  http://hepsoftwarefoundation.org/
- European Open Science Cloud Pilot Project launched …

  https://indico.cern.ch/event/319745/contribution/8/material/paper/0.pdf

Network in Switzerland

No issues !!!! Switch takes good care.

Swiss Network (opt)

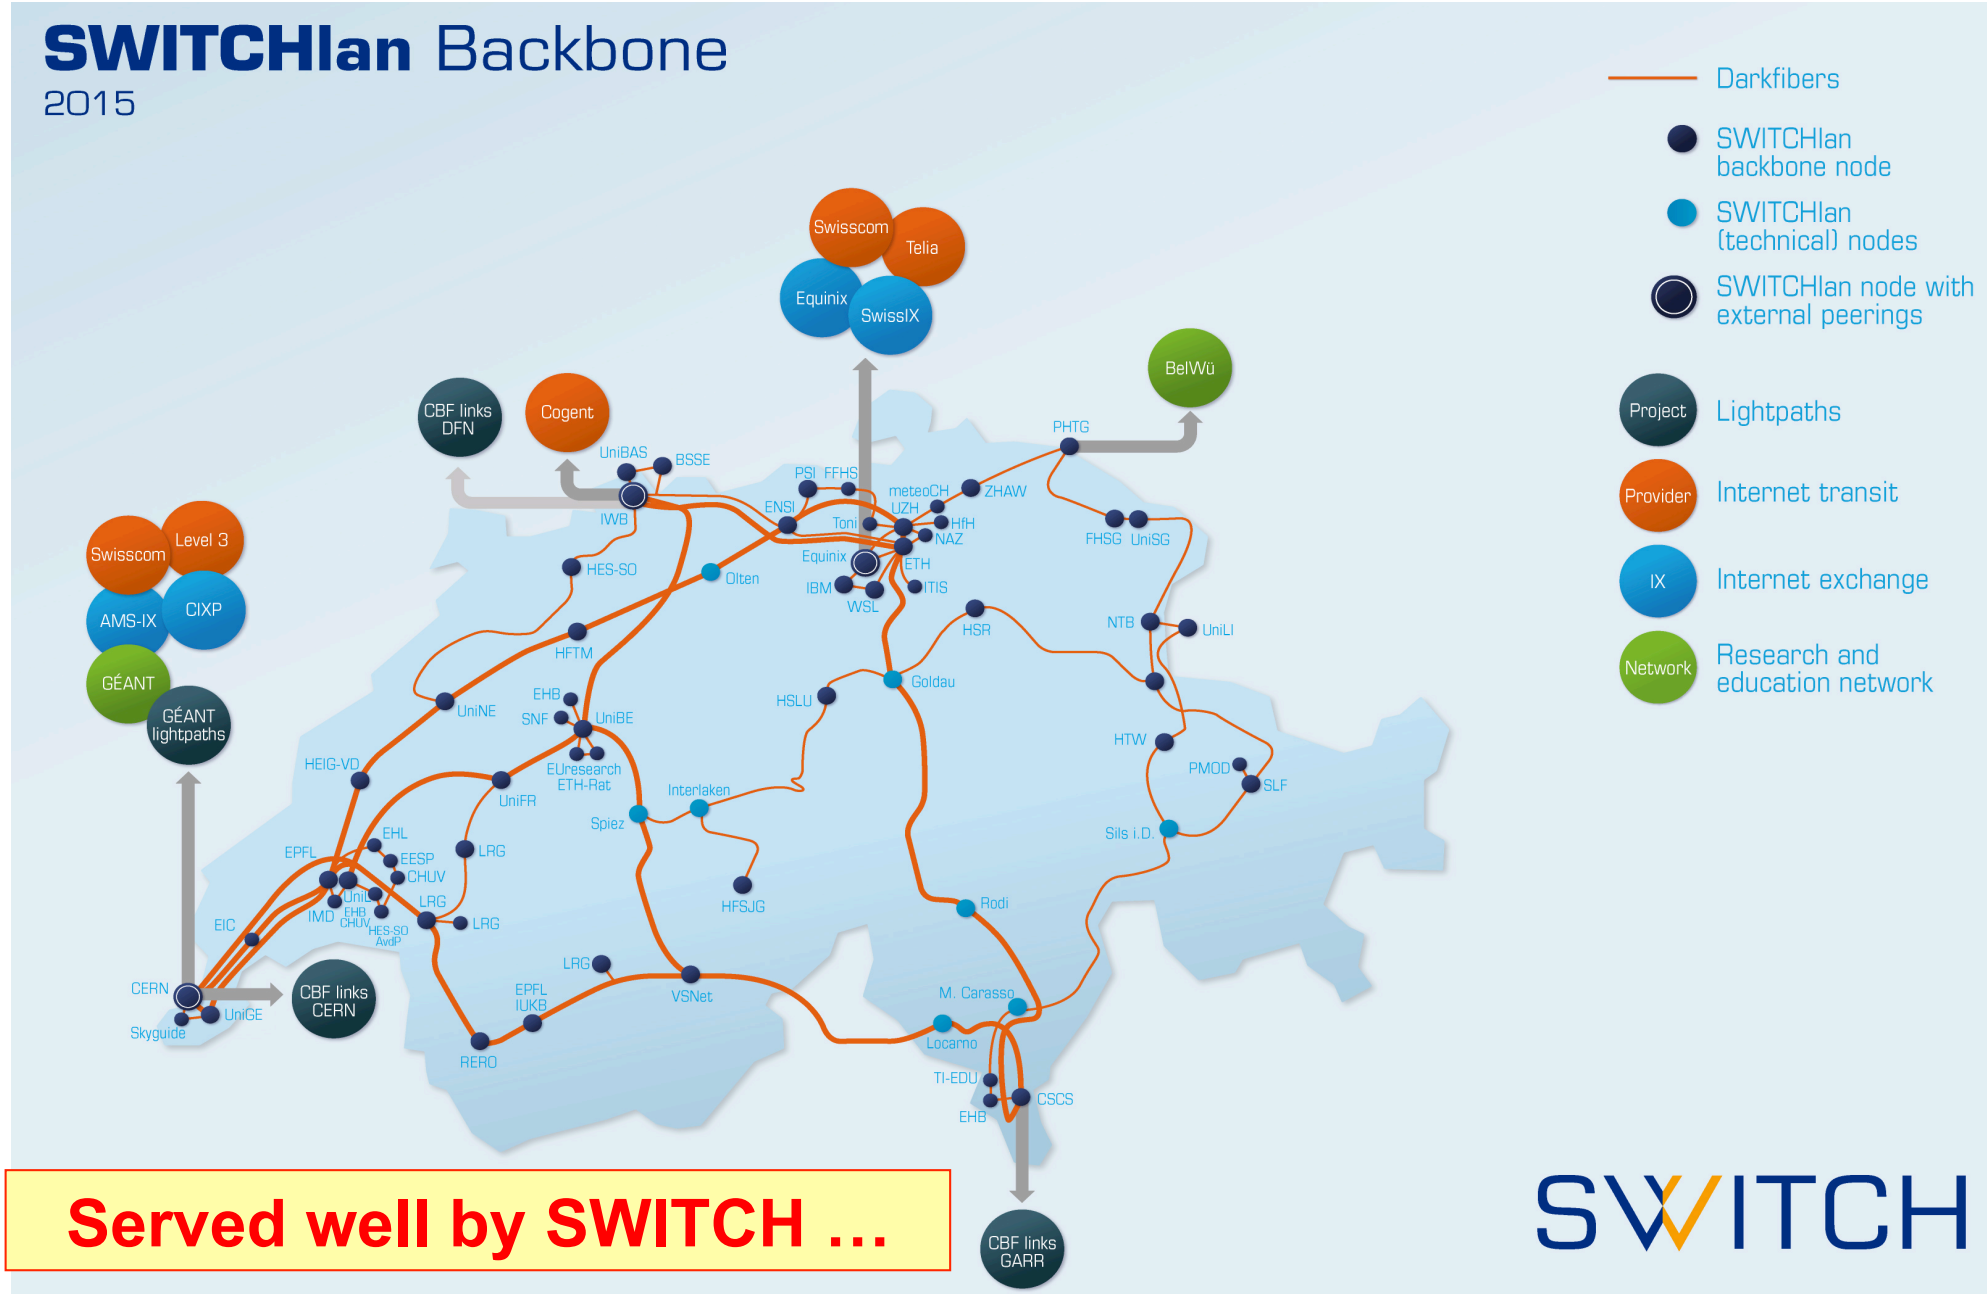

Served well by SWITCH …

SWITCH

WLCG Resource Usage Plots from CRRB 2015

See presentations at the C-RRB, April 2015
https://indico.cern.ch/event/359409/

Resources Usage– from CRRB

CPU history: efficiency

Experiments have become more efficient !

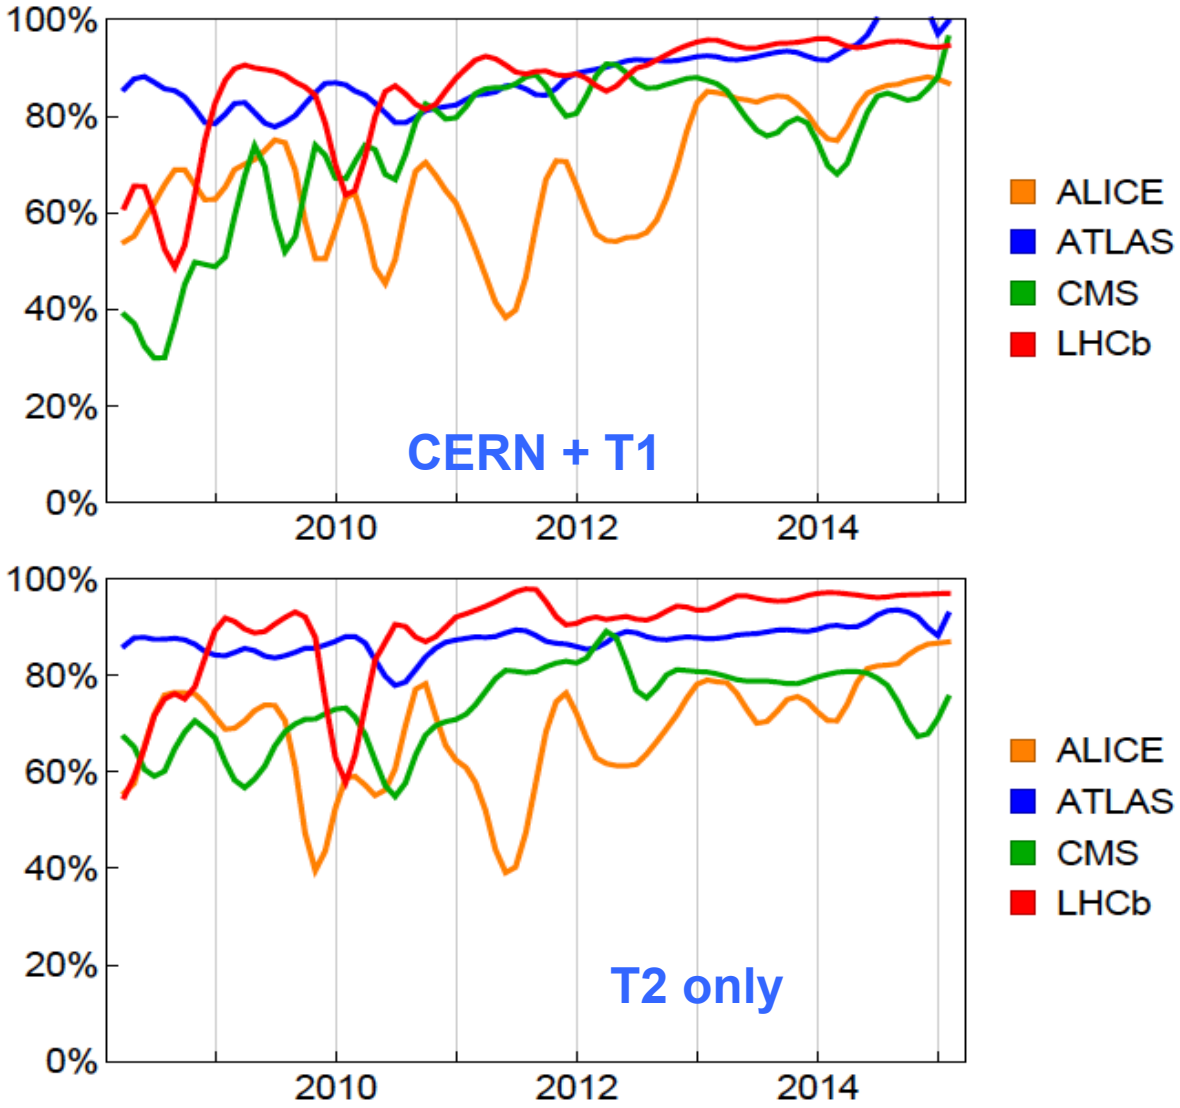

Top: CERN plus T1
Bottom: T2

Normalised CPU time (HS06·hrs) over normalised elapsed time.

Data from EGI accounting

Gaussian-smoothed monthly values

Data popularity: ATLAS accesses in time X

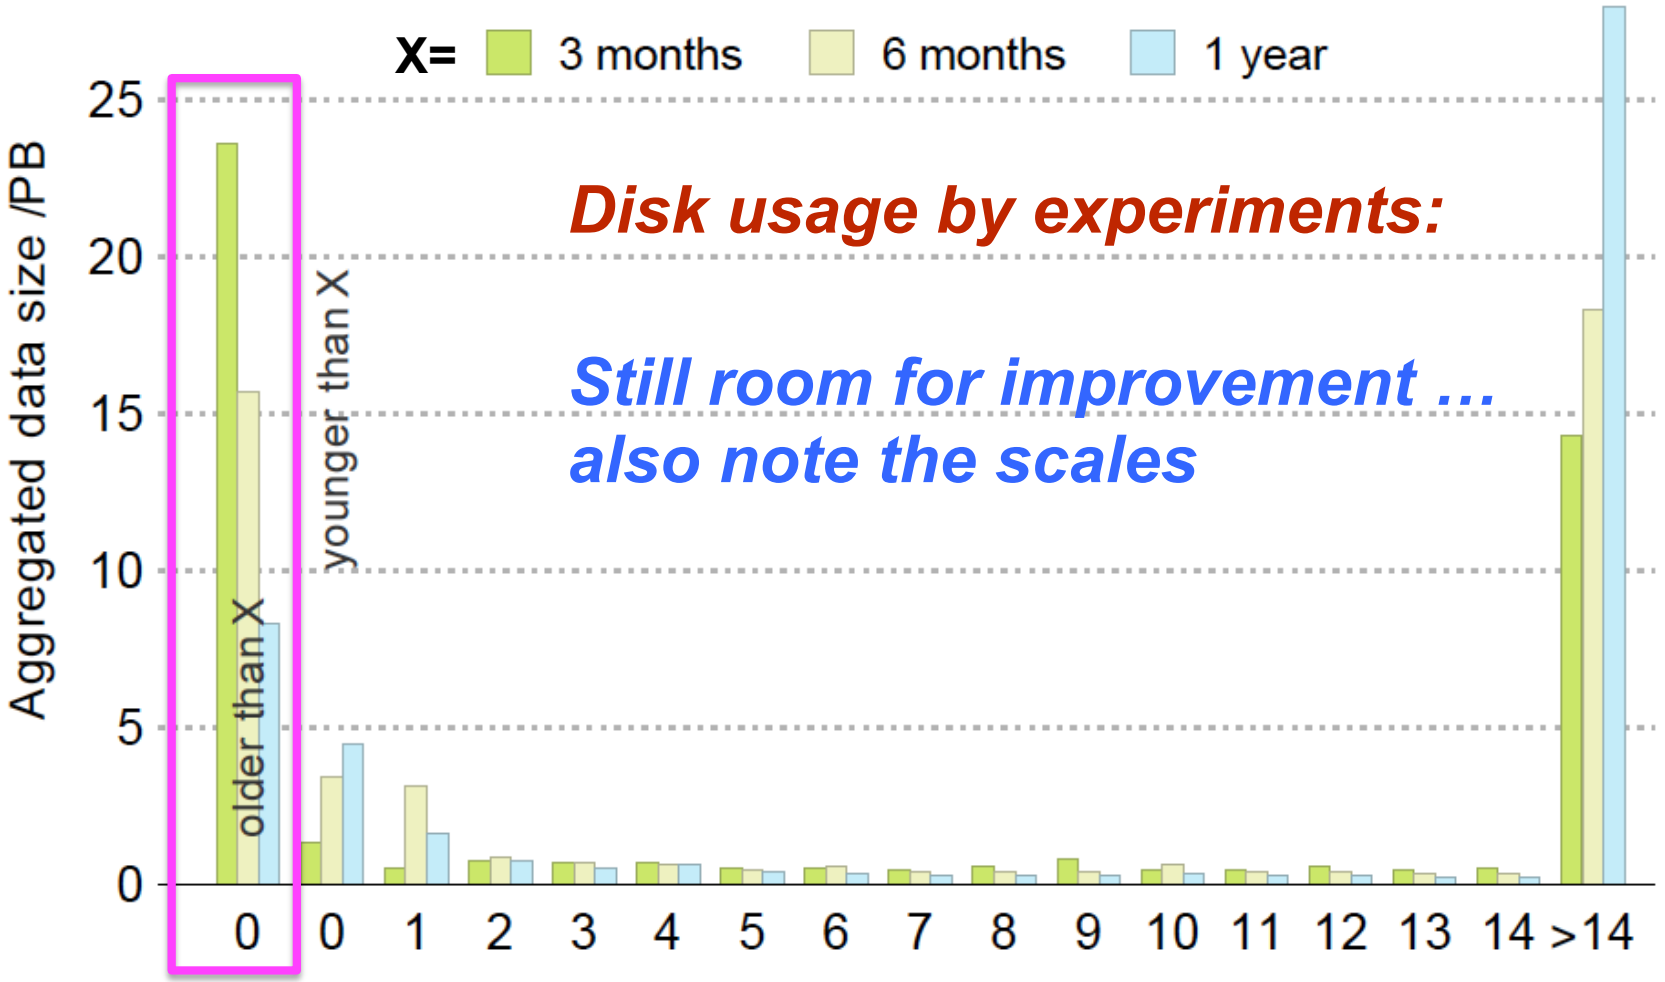

Volume of data versus number of accesses in ATLAS DATADISK at T1s and T2s for 3, 6 and 12 months to end of 2014

Data popularity: CMS accesses in time X

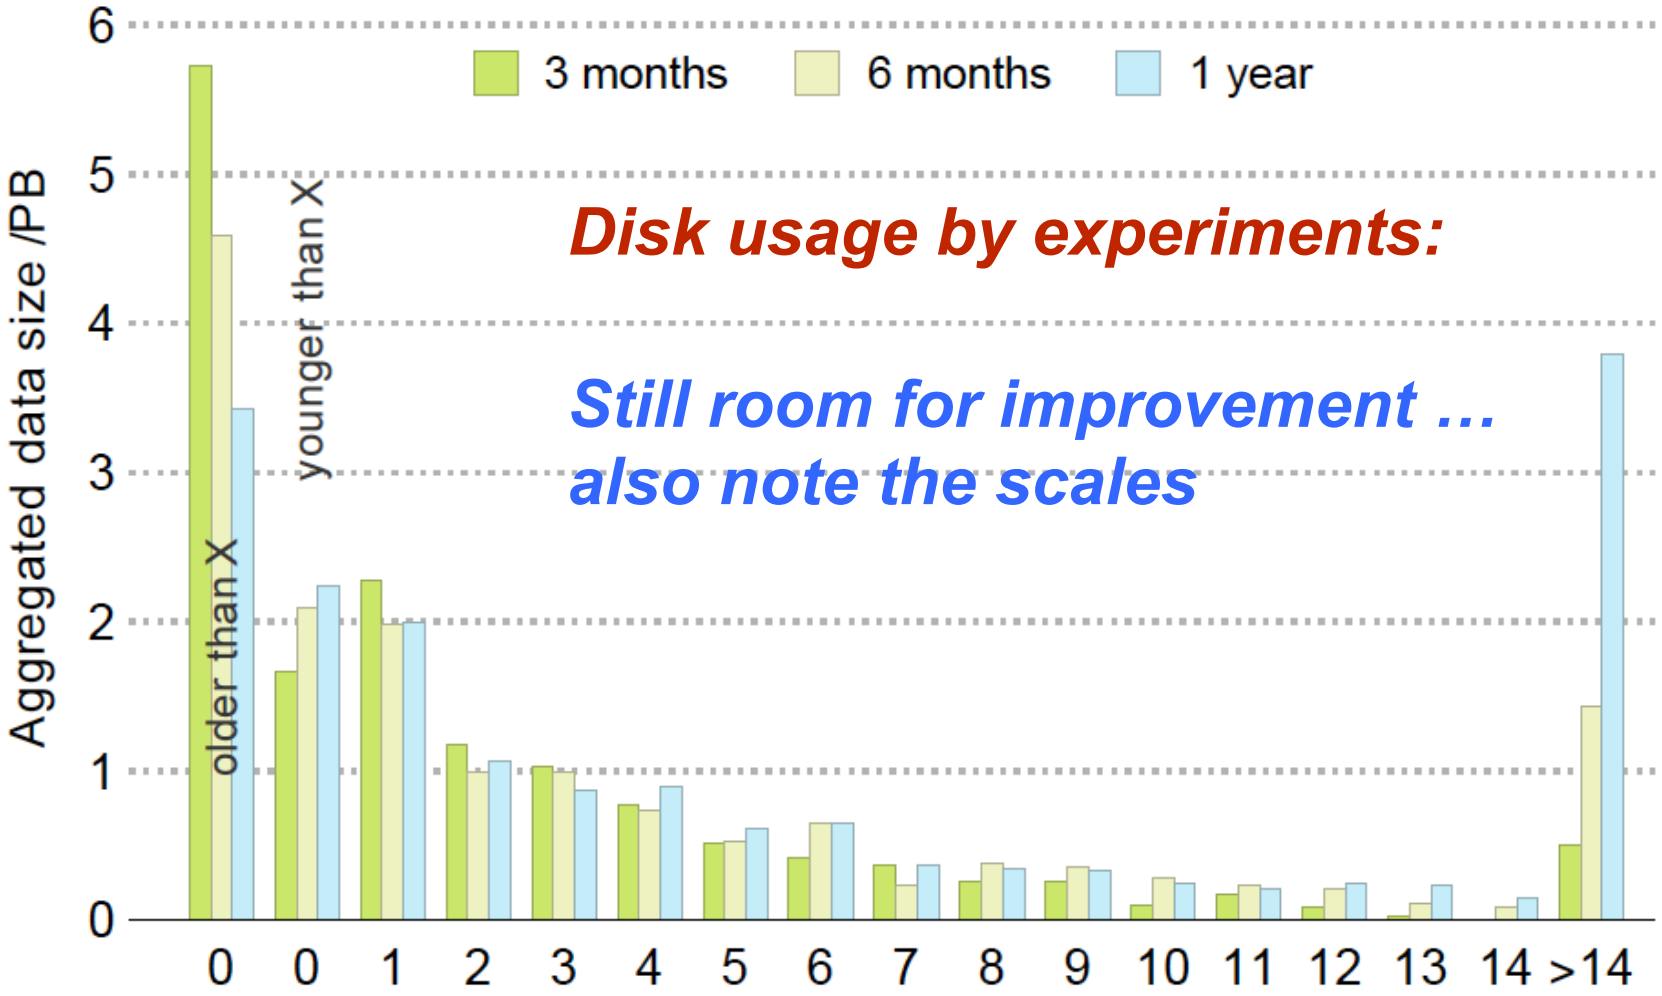

Volume of data versus number of accesses for CMS centrally-managed Tier 2 disk, for 3, 6 and 12 months

Data popularity: LHCb accesses in time X

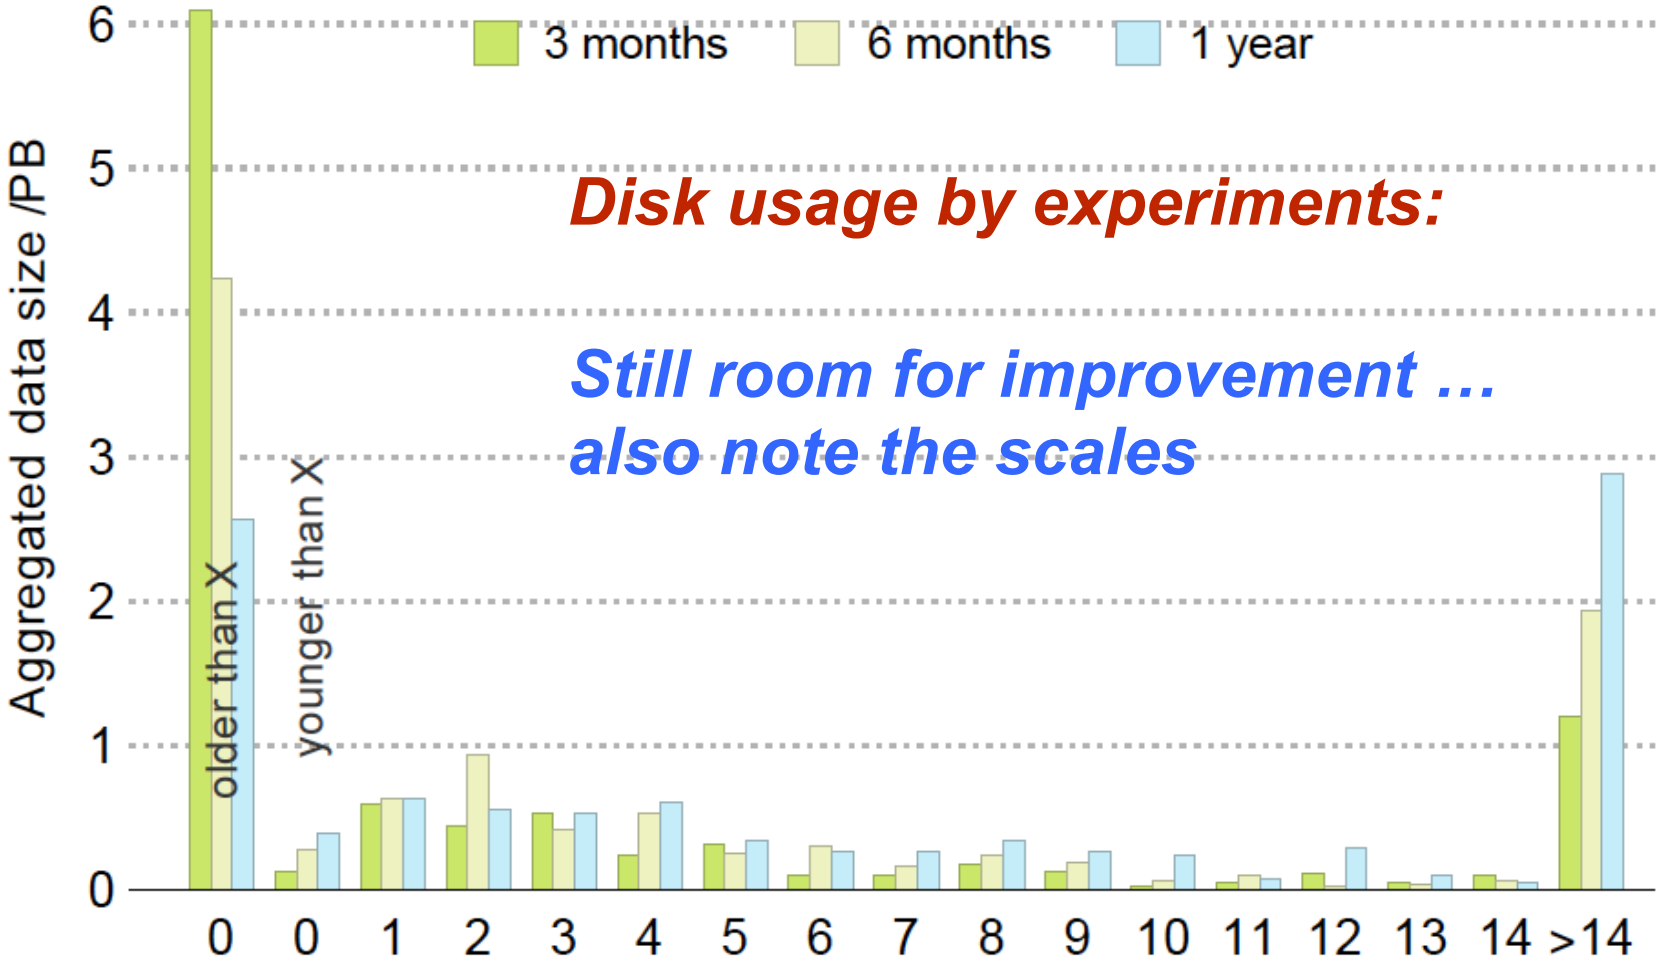

Volume of data versus number of accesses for LHCb for previous 3, 6 and 12 months

… We are ready for run 2

and are eagerly waiting for more data …

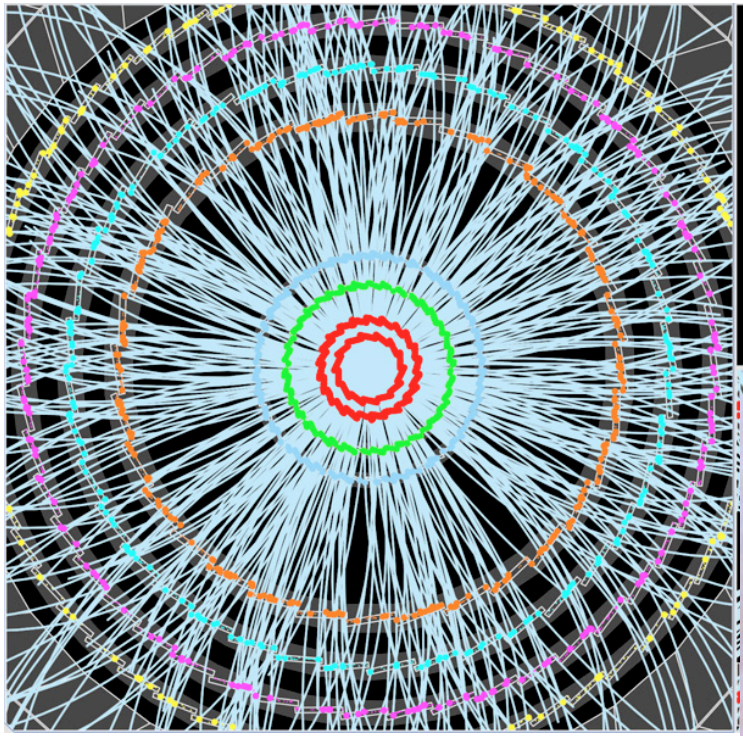

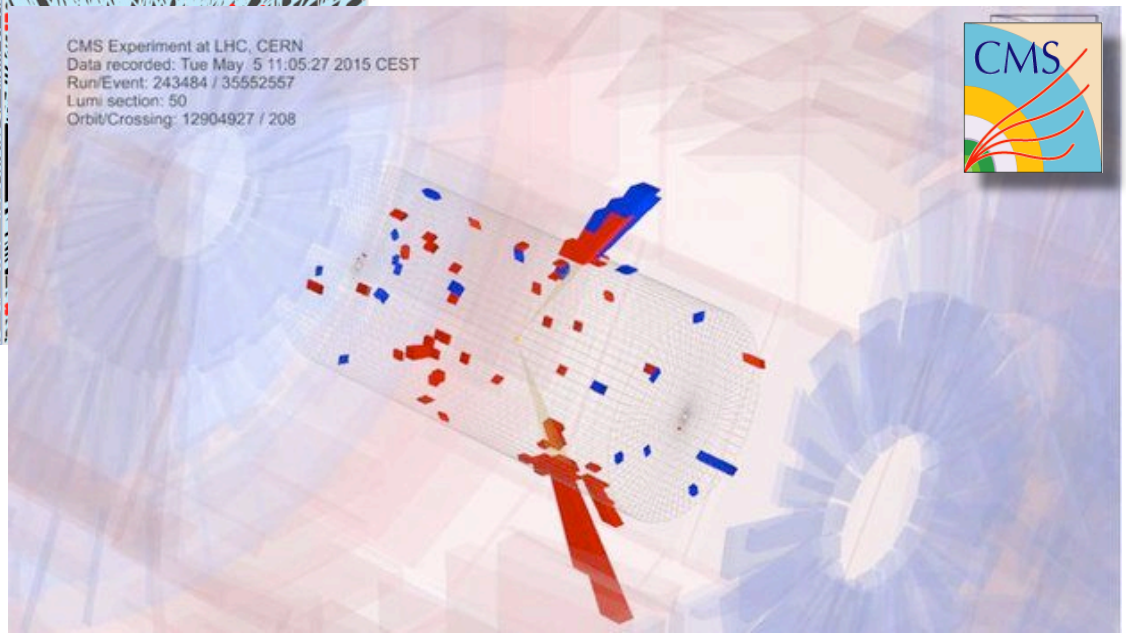

CHIPP Computing Board

Coordinates the Tier-2 activities
representatives of all institutions and experiments, CSCS, and included also the tier-3 expert

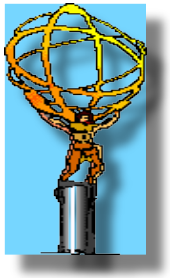

P.Iacobucci (UNI Ge)
S.Haug, G.Sciacca (UNI Bern)

C.Grab (ETHZ) chair CCB
D.Feichtinger (PSI) vice-chair CCB
D.Meister (ETHZ), F.Martinelli (PSI)

R.Bernet (UNIZH)
S.Tourneur (EPFL)

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

Thank you …

Backup slides

Resource comparisons – (opt)

WLCG Tier 2 Resources Situation on 20 April 2015							CERN-RRB-2015-012 Annex 2		

Switzerland, CHIPP, Manno	2014	2015	2016	Split 2015	ALICE	ATLAS	CMS	LHCb	SUM 2015
CPU (HEP-SPEC06)	37000	50000	50000	Offered	0	29000	14000	7000	50000
				% of Total		6%	3%	11%	5%
Disk (Tbytes)	2150	2350	2350	Offered	0	1225	875	250	2350
				% of Total		2%	3%	13%	3%

Estonia, NICPB, Tallinn	2014	2015	2016	Split 2015	ALICE	ATLAS	CMS	LHCb	SUM 2015
CPU (HEP-SPEC06)	45000	45000	45000	Offered	0	0	45000	0	45000
				% of Total			9%		9%
Disk (Tbytes)	1000	1000	1000	Offered	0	0	1000	0	1000
				% of Total			3%		3%

Romania, Romanian Tier-2 Federation	2014	2015	2016	Split 2015	ALICE	ATLAS	CMS	LHCb	SUM 2015
CPU (HEP-SPEC06)	38000	41600	41600	Offered	17000	19700	0	4900	41600
				% of Total	9%	4%		7%	5%
Disk (Tbytes)	2330	2690	2690	Offered	1280	1087	0	323	2690
				% of Total	6%	2%	0%	17%	3%

CERN Tier0 / CAF	2014	2015	2016	Split 2015	ALICE	ATLAS	CMS	LHCb	SUM 2015
CPU (HEP-SPEC06)	356'000	687'000	840'000	Offered	175000	205000	271000	36000	687000
				Required	175000	205000	271000	36000	687000
				% of Req.	100%	100%	100%	100%	100%
Disk (Tbytes)	29'100	49'000	57'500	Offered	14500	14000	15000	5500	49000
				Required	14500	14000	15000	5500	49000
				% of Req.	100%	100%	100%	100%	100%
Tape (Tbytes)	82'000	95'400	128'200	Offered	16200	33000	35000	11200	95400
				Required	16200	33000	35000	11200	95400
				% of Req.	100%	100%	100%	100%	100%

Swiss Tier-2 CPU Planning (opt)

Split CPU (kHS06) per VO of resources pledged to WLCG.

Experiment	2012 (delivered)	2013 (delivered) Phase G	2014 (delivered) Phase H	2015 (projected) Phase J	2016 (future) Phase K	2017 (future) Phase L
ATLAS CSCS AEC-UNIBE	7.0	9.2 5	10.4 11	14 11	14 11 ?	15 ?
CMS	7.0	9.2	10.4	14	14	15
LHCb	3.5	4.6	5.2	7	10	14
TOTAL (kHS06)	**17.4**	**23**	**26**	**35**	**39**	**44**

Note: the ATLAS numbers list CSCS and AEC-UNIBE numbers separately.
AEC-UNIBE contributes to ATLAS only.

Swiss Tier-2 Disk Planning (opt)

Split storage per VO of resources pledged to WLCG

Experiment	2012 (delivered)	2013 (delivered) Phase G	2014 (delivered) Phase H	2015 (projected) Phase J	2016 (future) Phase K	2017 (future) Phase L
ATLAS CSCS AEC-UNIBE	544	649	792 350	875 350	955 350 ?	1040 ?
CMS	544	649	792	875	955	1040
LHCb	2	2	216	550	690	820
TOTAL (TB)	**1090**	**1300**	**1800**	**2300**	**2600**	**2900**

Note: the ATLAS numbers list CSCS and AEC-UNIBE numbers separately.
AEC-UNIBE contributes to ATLAS only.

Storage usage at CH-Tier2

UPDATED

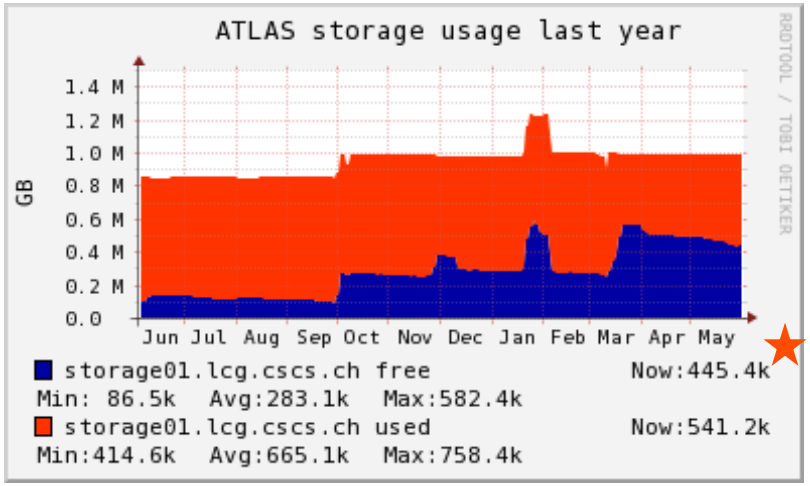

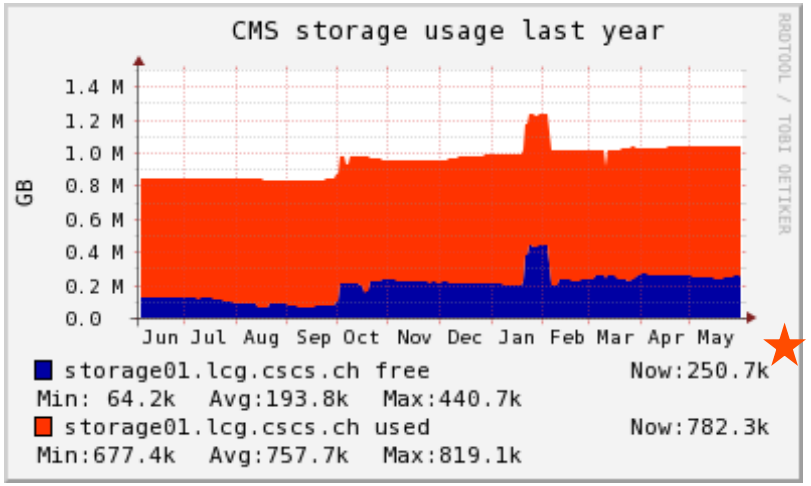

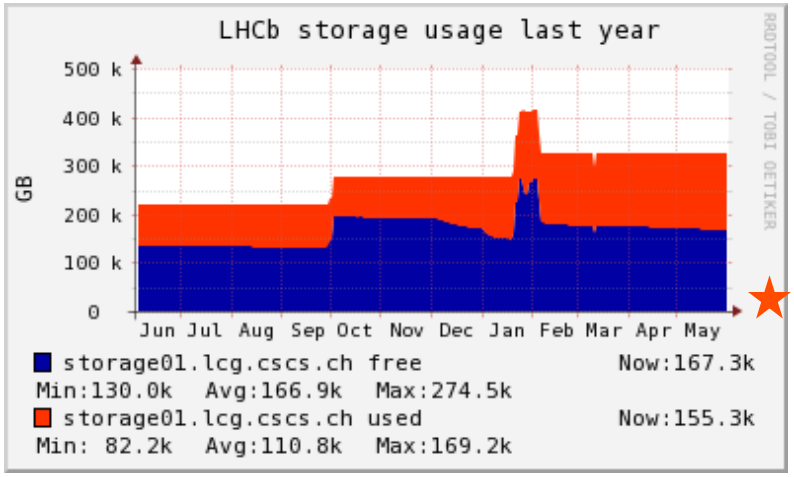

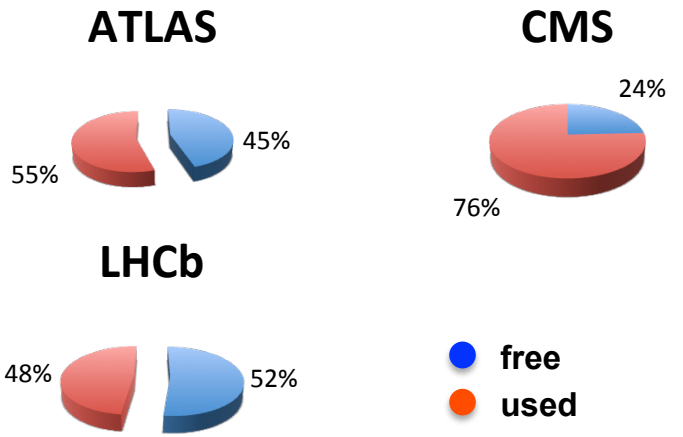

★ Values expressed in Gigabyte as shown by the filesystem

- Overall good exploitation
- LHCb started to exploit obtained storage …

http://ganglia.lcg.cscs.ch/ganglia3/?r=year&cs=&ce=&tab=v&vn=DCACHE+storage+distribution

Comment on Certificates

SWITCH

SWITCHpki Features Participants Sign up Certificate Management Documents Contact

Certificate renewal/replacement

Notice – April 2015

SWITCH will stop issuing and renewing Grid certificates on 30 September 2015, as announced in a message to all holders of valid certificates on 9 March 2015, and all Grid certificates issued since end of March 2015 will only be valid through 31 March 2016. Users are encouraged to get in touch with SwiNG, the Swiss National Grid Association, to get more information about their future Grid certificate offering.

European Open Science Cloud Pilot Project

- Bring together the stakeholders
  - Research Infrastructures *(ESFRI, etc.)*
  - Research Organisations (WLCG tier-1 etc.)
  - European e-Infrastructures *(GEANT, EGI, PRACE, EUDAT, OpenAIRE)*
  - Commercial cloud service providers *(Helix Nebula, etc.)*
  - End-users including the *long-tail of science*

- Deliver the pilot
  - Technical architecture for the hybrid cloud
  - Security model compatible with EU data protection legislation
  - Assemble and deploy a 5% scale prototype
  - Verify the business model to ensure it can be sustained beyond the pilot
  - Governance structure avoiding monopoly of any research group or service provider
  - Roadmap for full-scale implementation