Spotkanie WLCG-PL i eksperymenty przy LHC

Michał Bluj

27 czerwca 2023

Plan wystąpenia

- Podsumowanie ostatniego okresu działalności
  - Informacja z ostatniego spotkania C-RRB (kwiecień 2023)
- Zobowiązania na 2023 i wstępne plany na 2024
- Dyskusja i AOB

Informacje ze spotkania C-RRB

- Zebranie odbyło się 25 IV 2023

14:00 → 14:02 **Introduction** — 2m

Speaker: Joachim Mnich

14:02 → 14:05 **Approval of the minutes of the last meeting** — 3m

CERN-RRB-2022-122

14:05 → 14:35 **Status of the WLCG project** — 30m

Speaker: Simone Campana

CERN-RRB-2023-048 | CERN-RRB-2023-04...

14:35 → 14:40 **LHCC Deliberations** — 5m

Speaker: Lorenzo Moneta (Scientific Secretary, LHCC)

CERN-RRB-2023-05... | Slides

14:40 → 15:00 **Computing Resources Scrutiny Group Report** — 20m

Speaker: (Remotely) Pekka Sinervo (Chairperson, CRSG)

CERN-RRB-2023-012 | CERN-RRB-2023_01... | CERN-RRB-2023_01...

15:00 → 15:05 **Summary** — 5m

Speaker: Joachim Mnich

WLCG news

- Za raportem projektu WLCG z C-RRB (S. Campana)

Two currently T2 sites (NCBJ Swierk – PL and IHEP – CN) expressed their intention to offer T1 services, initially to LHCb at the level of 5% T1 capacity each

- The proposals were discussed at the December 2022 WLCG Overview Board. The two proposals were endorsed.
- NCBJ Swierk and IHEP are now protoT1 sites and will be commissioned as T1s in the next year, in direct collaboration with LHCb
- The WLCG Management Board is organising and monitoring this process. An update is expected at the June 2023 WLCG Overview Board

ASGC steps down as ATLAS T1 in October 2023

WLCG started tackling cases where a Federation signed an MoU but has been inactive for many years. The process was discussed at the WLCG Overview Board

- One federation (KAVALA - GR) was contacted and asked about the future intentions. The federation notified WLCG that it does not intend to continue providing services
- KAVALA's participation in WLCG has been SUSPENDED and the MoU should be considered not effective any longer

Koszty energii

- Za raportem projektu WLCG z C-RRB (S. Campana)

Different initiatives to face the increasing cost of energy

- Reduction of capacity in critical months or for the full year
- (R&D) Real time power load shedding to compensate peaks of energy demand

Germany: load shedding based on energy cost

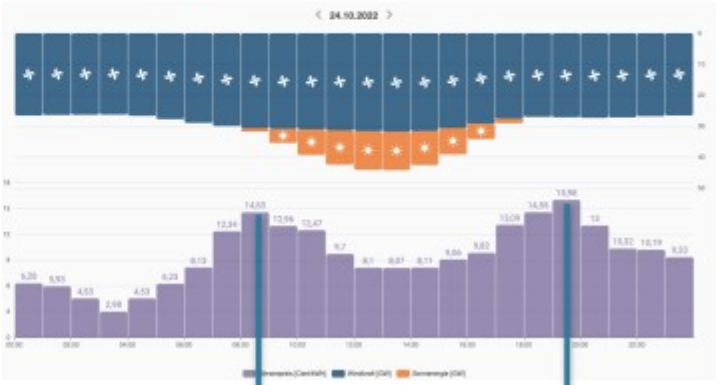

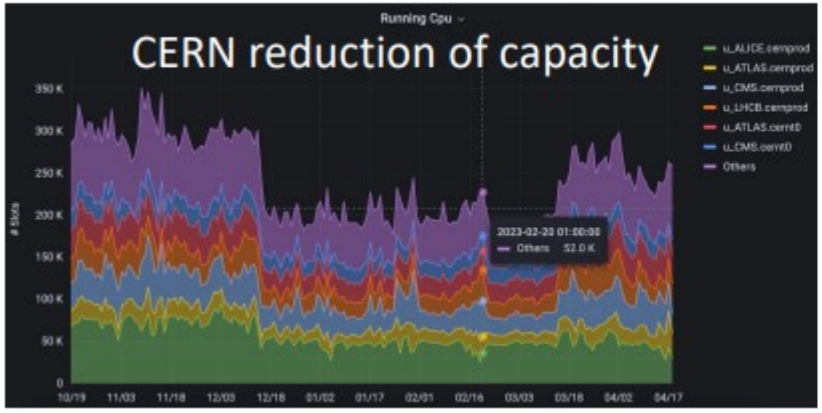

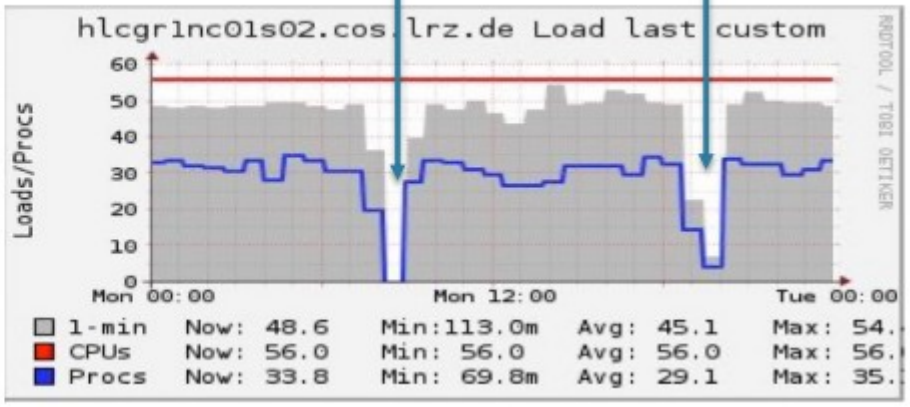

Flat budget model

- Za raportem projektu WLCG z C-RRB (S. Campana)

The WLCG "flat budget model" agreed with the RRB: **with the same level of funding for computing hardware every year we expect to increase the amount of resources by 15% every year**

We started a study of the purchase power variation trend over the years (CERN and UK). Next steps will be enlarging to other countries and refine the model

CPU average variation (5 years): +22%

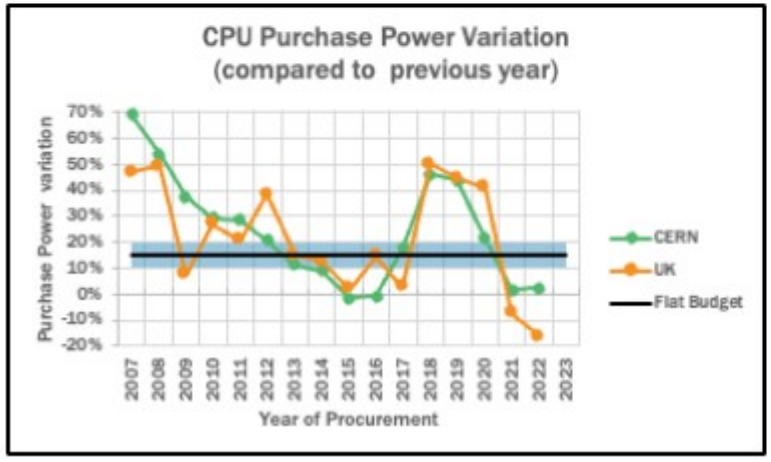

DISK average variation (5 years): +13%

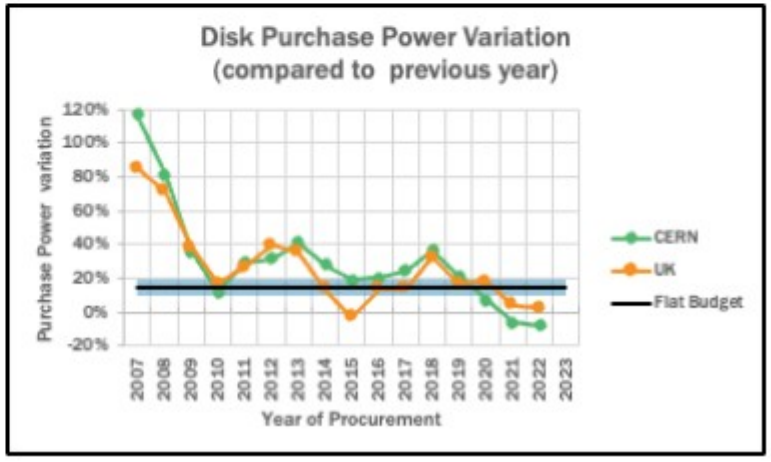

- Wygląda, że po pandemii trend "+15%" powinien zostać utrzymany
  - Ale aby się utrzymać trzeba być gotowym na zmianę architektury x86->ARM64
- Wygląda, że wzrost „+15%" pokrywa wymagania eksperymentów (nestępny slajd)

Flat budget model

- Za raportem projektu WLCG z C-RRB (S. Campana)

We used the findings from the study on the previous slide to understand

- The difference between the “15% Flat Budget model” (more theoretical) and a “Measured Flat Budget model” (based on the real purchase power)
- How the WLCG resource needs compare to the two models

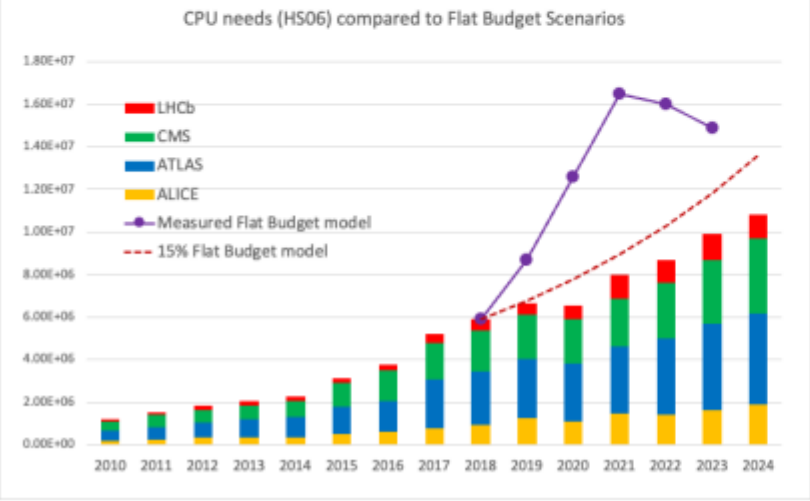

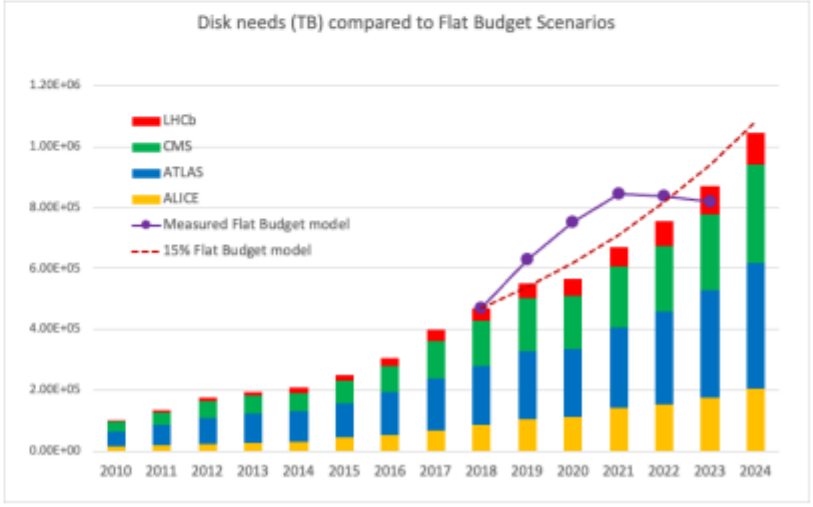

- Wygląda, że wzrost „+15%” pokrywa wymagania eksperymentów

Architektury inne niż x86 (GPU)

- Za raportem projektu WLCG z C-RRB (S. Campana)

GPUs for offline processing

CMS High Level Trigger: CPU + GPU (Nvidia T4, 2 per node) solution. CMSSW data processing framework already capable to orchestrate computations to accelerators

- 40% of evt. processing offloaded, +70% throughput, +50% per kWh

CMS is working to profit from accelerator-ready code for offline processing already in Run 3. Solid plan ahead: aim to 10% offload this year

ALICE is running asynchronous reconstruction productions with CPU+GPU in O2 since January 2023. This allows to:

- Measure the speed up in processing
- Understand the challenges in job tuning in view GPU use on the Grid.

We are getting closer to be able to benefit from GPUs in WLCG

Architektury inne niż x86

- Za raportem projektu WLCG z C-RRB (S. Campana)

It is not just about GPUs but also non-X86 CPUs

- Power CPUs: CMS software validated in 2022 on the INFN M100 HPC at CINECA during 2022. Machine now used in production

- ARM64: most of the LHC workloads ported to ARM. ATLAS software successfully validated. Promising results in terms of energy per processed event

Give options to WLCG sites to diversify their CPU architectures, based on overall cost of ownership considerations

IBM Power CPU architecture

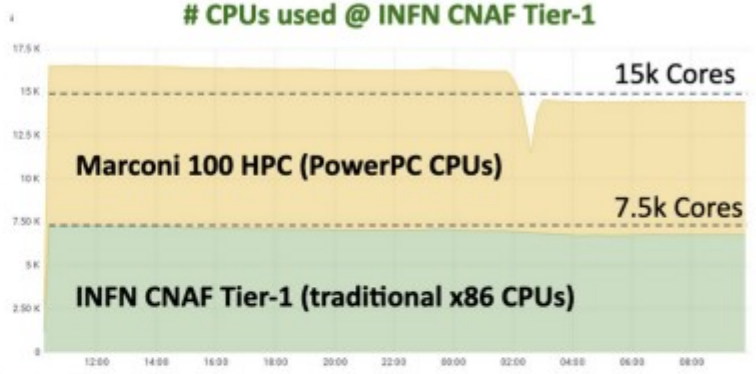

- Measured energy-used per unit WL-Score for 2 ATLAS and 3 CMS HEP-Score workloads.
- Three test machines:
  - x86_64: Single AMD EPYC 7003 series (SuperMicro)
  - 2*x86_64: Dual AMD EPYC 7513 series Processors (DELL)
  - arm64: Single socket Ampere Altra Processor (SuperMicro)
- Normalised to (first) x86. Lower is better!

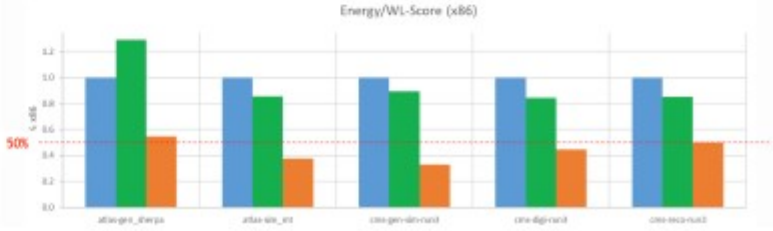

Oczekiwania na 2024

- Oczekiwania eksperymentów na 2024 (za raportem Computing Resources Scrutiny Group z C-RRB)
  - Pledges w IX'23, (dostępność od IV’24 do III’25)

- ”Big picture” view for 2023
  - Continuing data collection and processing
  - Increased focus on Run 3 physics analyses
  - ALICE expected to acquire large dataset
  - LHCb making best use of 2023 detector configuration
- “Big picture” view for 2024
  - Increases ~ 10% for CPU, ~18% for disk, ~25% for tape
  - ALICE, ATLAS and CMS request large increases in disk and tape
    - Reflects Run 3 data-taking and physics analysis
  - LHCb expects no change from 2023 levels

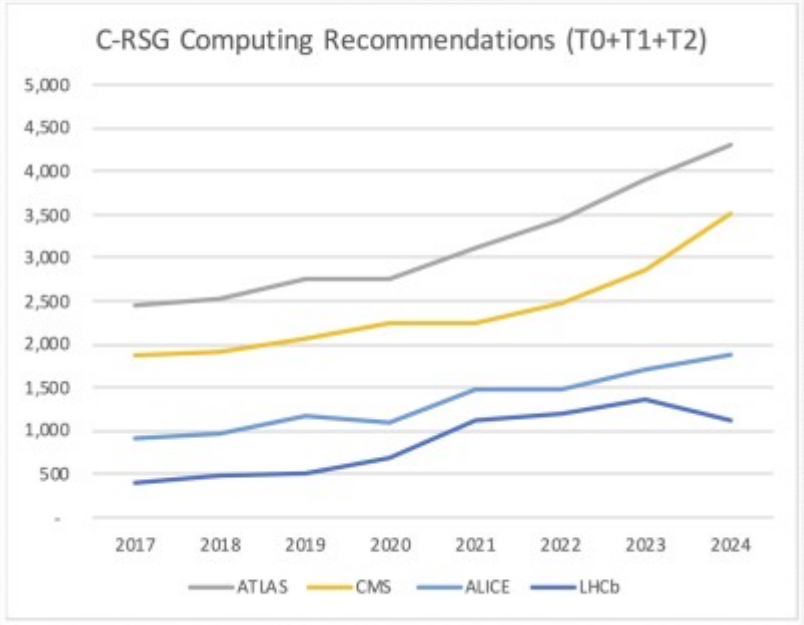

Pekka Sinervo, C.M.

April 24, 2023

Oczekiwania na 2024

- Oczekiwania eksperymentów na 2024 (za raportem Computing Resources Scrutiny Group z C-RRB)
  - Pledges w IX'23, (dostępność od IV'24 do III'25)

ALICE

		2024		
		Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	600	111%	600
	Tier-1	630	110%	630
	Tier-2	650	110%	650
	HLT	n/a	n/a	n/a
	Total	**1880**	**110%**	**1880**
	Others			
Disk	Tier-0	67.5	115%	67.5
	Tier-1	71.5	113%	71.5
	Tier-2	66.5	116%	66.5
	Total	**205.5**	**114%**	**205.5**
Tape	Tier-0	181.0	138%	181.0
	Tier-1	107.0	130%	107.0
	Total	**288.0**	**135%**	**288.0**

- 2022 and 2023 priorities
  - Analyze 2022 pp data
  - Take and analyze larger PbPb data sample in 2023
- 2024 reflects data set growth
  - Plans for pp, PbPb and pPb collisions
- Will still be working to implement more aggressive compression
- Increase in MC/data by x4

ATLAS

		2024		
		Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	936	126%	936
	Tier-1	1516	106%	1516
	Tier-2	1852	106%	1852
	HLT	n/a	n/a	n/a
	Total	**4304**	**110%**	**4304**
Disk	Tier-0	49.0	123%	49.0
	Tier-1	163.0	120%	163.0
	Tier-2	200.0	119%	200.0
	Total	**412.0**	**120%**	**412.0**
Tape	Tier-0	207.0	119%	207.0
	Tier-1	452.0	128%	452.0
	Total	**659.0**	**125%**	**659.0**

- 2022 and 2023 has Run 3 focus
  - Higher pileup -> 15% CPU increase
  - Moving to smaller data formats
- 2024 requests reflect full year of running to collect 110 fb^{-1} of data
  - CPU increase reflects improved simulation performance
  - Doesn't account for HI run -- OK
- Porting to other architectures
  - Both ARM and GPUs
  - Continued use of opportunistic resources
- HL-LHC work remains priority

CMS

		2024		
		Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	980	136%	980
	Tier-1	930	116%	930
	Tier-2	1600	119%	1600
	HLT	n/a	n/a	n/a
	Total	**3510**	**122%**	**3510**
Disk	Tier-0	54.0	120%	54.0
	Tier-1	122.0	124%	122.0
	Tier-2	149.0	127%	149.0
	Total	**325.0**	**125%**	**325.0**
Tape	Tier-0	320.0	140%	320.0
	Tier-1	380.0	120%	380.0
	Total	**700.0**	**129%**	**700.0**

- 2022 and 2023 focused on Run 3
  - Increased T0 resources due to higher pileup
  - Continued migration to nano-DST
- 2024 requests reflect higher pileup and larger dataset
- T0 increase result of larger data set and higher
- Increased effort on HL-LHC simulation

LHCb

		2024		
		Request	2024 req. /2023 C-RSG	C-RSG recomm.
CPU	Tier-0	174	81%	174
	Tier-1	572	81%	572
	Tier-2	319	82%	319
	HLT	50	n/a	50
	Total	**1115**	**82%**	**1115**
Disk	Tier-0	30.6	101%	30.6
	Tier-1	61.2	101%	61.2
	Tier-2	11.8	102%	11.8
	Total	**103.6**	**101%**	**103.6**
Tape	Tier-0	117.1	129%	117.1
	Tier-1	133.3	85%	133.3
	Total	**250.4**	**101%**	**250.4**

- Simulation dominates LHCb 2022 and 2023 CPU usage
  - Disk needed for data sprucing and analysis
  - Reduced data set in 2023
- 2024 request reduced due to status of detector
  - Data used for commissioning and some physics studies
- Working to improve simulation in high-pileup environment
- Developing new T1 sites

WLCG poza LHC

- Za raportem projektu WLCG z C-RRB (S. Campana)

Many communities collaborate with WLCG, share some of the same technologies, services and resources

Fostering this collaboration is essential for the sustainability of the WLCG infrastructure. See the WLCG/DUNE/Belle-2 Snowmass contribution

DUNE, Belle-2, JUNO **and (recently) VIRGO** are now WLCG "observers"

ECFA, NuPECC and APPEC started to jointly discuss a strategy and a possible implementation of European federated computing at future large-scale research facilities, as recommended by the funding agencies

- A first workshop will happen in June 2023

- WLCG sees this also as an opportunity to prepare a common position for the next European Strategy for Particle Physics

Posumowanie

- Dobre działanie WLCG-PL w 2022 i w 1 poł. 2023
- Oczekiwania eksperymentów na 2024
  - "Adiabatyczne" zwiększenie zasobów ALICE, ATLAS, CMS (+10-15%)
  - Przygotowanie T1 dla LHCb
- Potencjalne problemy „systemowe” dostrzegane przez WLCG (bez zmian)
  - Problemy z cenami energii (dotyczy PL),
  - Problem z cenami i czasem dostaw sprzętu
- Przygotowanie do obliczń z użyciem niehomogenicznych architektur
  - GPU, ARM64