Meeting annuale ATLAS-IT calcolo : introduzione

(L. Carminati, A. Doria, L.Rinaldi)

Genova, Novembre 2023

Intro

- Benvenuti a tutti ! (e un grazie speciale a Dario e Mirko per l'ospitalità)
- Agenda “flessibile” : prendiamoci il tempo necessario per discutere per bene tutti gli argomenti
    - Martedì : intervento ATLAS S&C coordinators, operations, stato dei siti, piani futuri
    - Mercoledì : spazio alle attività di S&C svolte da membri di ATLAS Italia
- Cena ore 20 : cerchiamo di chiudere entro le 18 in modo da avere il tempo di passare in albergo prima della cena

Outline

- [Overview delle risorse di computing utilizzate dall’esperimento ATLAS]
- Computing ATLAS svolto sulle risorse Italiane
- Stato e funzionamento delle nostre risorse : performance, pledge
- Crescita computing ATLAS Italia per il 2024
- Aree di sviluppo e challenges per il computing ad HL-LHC

Attività ATLAS al T1 : CPU (settembre 2022-settembre 2023)

- Pledge 121 kHS23
- Average CNAF = 115 kHS23

Recente tuning degli algoritmi di sharing e priorità'

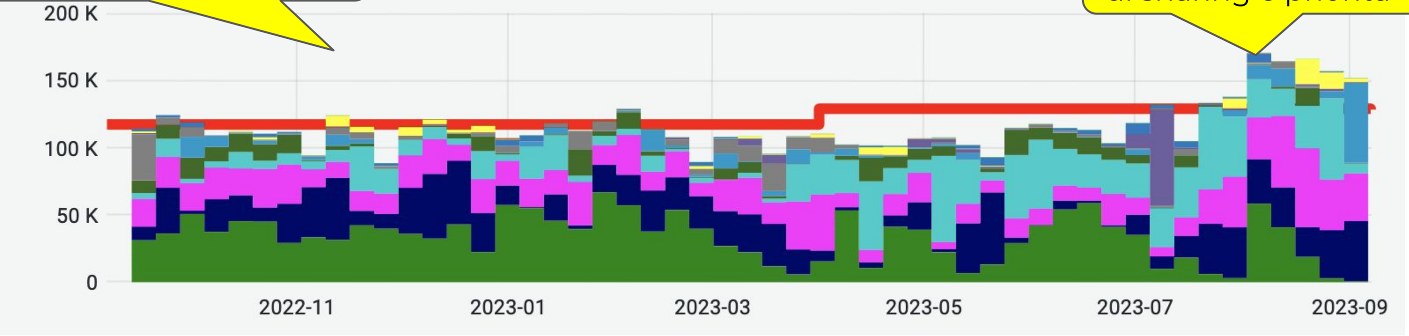

- ❑ Breakdown attività
    - ❑ produzione di MC (generazione, simulazione e ricostruzione)
    - ❑ preparazione dati secondari per l'analisi (group production),
    - ❑ analisi (user analysis)

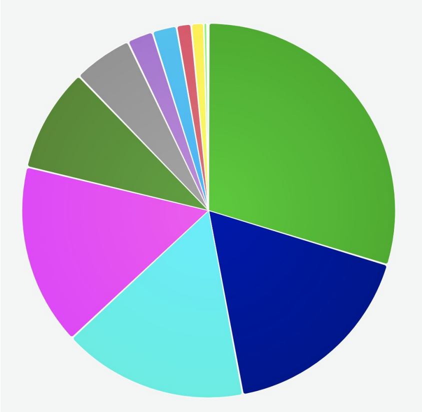

	Value	Percent
Group Production	1.02 Tri	30%
MC Reconstruction	595 Bil	17%
MC Simulation Full	548 Bil	16%
User Analysis	537 Bil	16%
MC Event Generation	307 Bil	9%
Group Analysis	173 Bil	5%
MC Merge	74.7 Bil	2%
MC Simulation Fast	69.9 Bil	2%
MC Resimulation	40.8 Bil	1%

Attività ATLAS al T1 : CPU (novembre 2022-novembre 2023)

- Pledge 125 kHS23
- Average CNAF = 128 kHS23

Con l'incremento di share della seconda parte dell'anno siamo in pari !

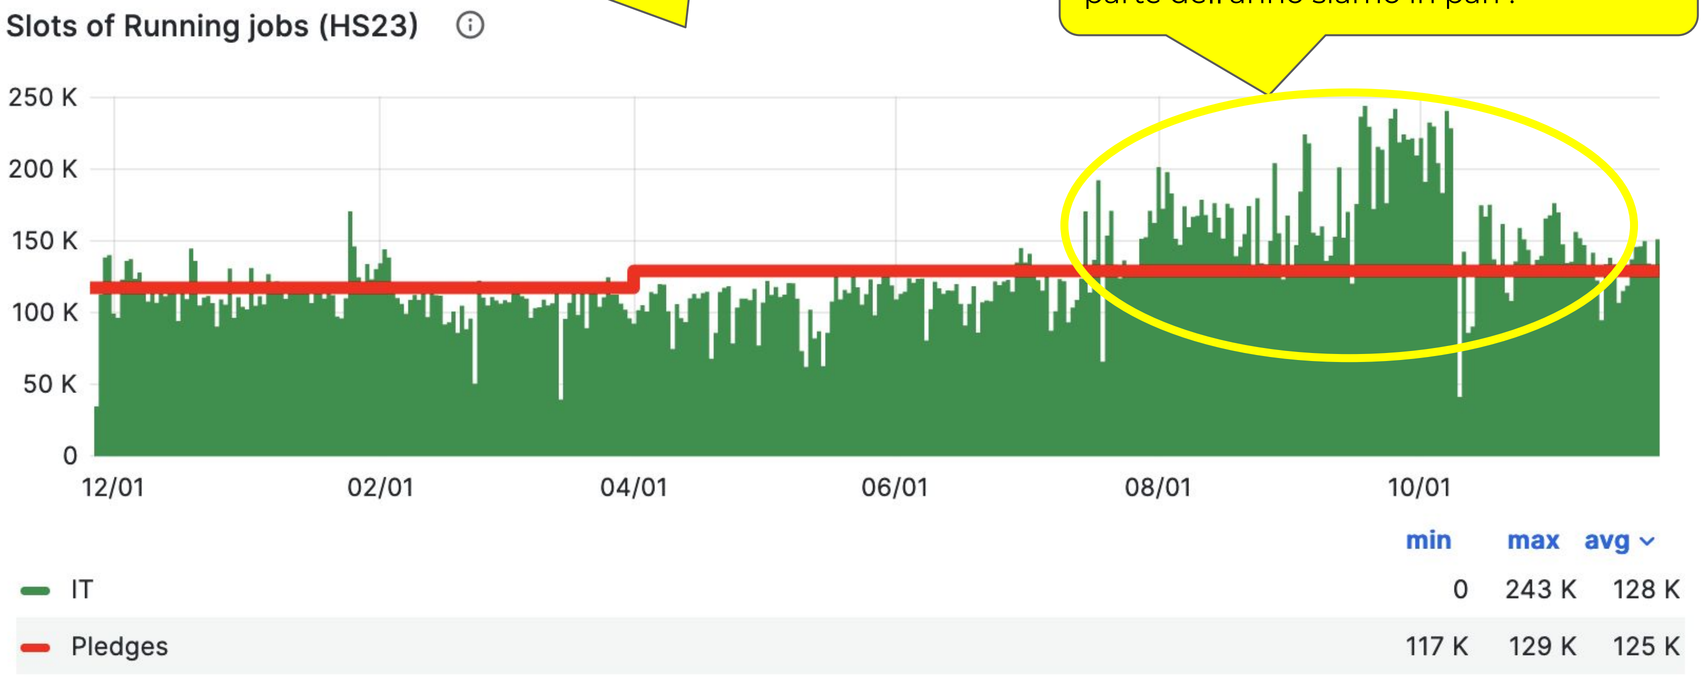

Attività al T1 : occupazione disco e tape

Referaggio 2023

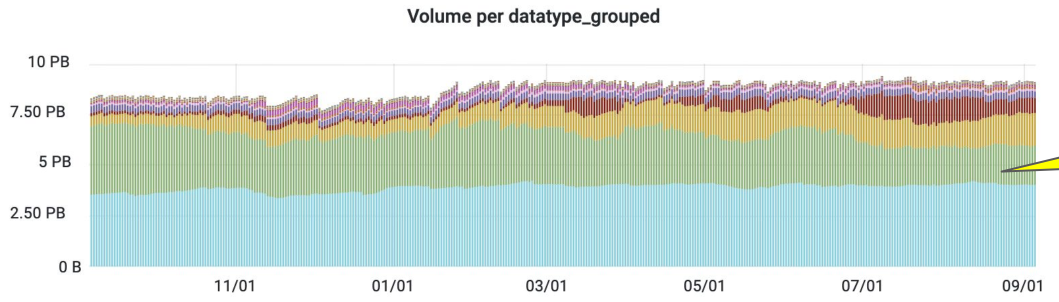

	min	max	avg	current
DAOD	0 B	4.24 PB	3.90 PB	4.04 PB
AOD	0 B	3.53 PB	2.63 PB	1.95 PB
HITS	0 B	1.64 PB	1.03 PB	1.61 PB

Disk

9.50 PB occupato (10.21 PB disponibile)

Tape

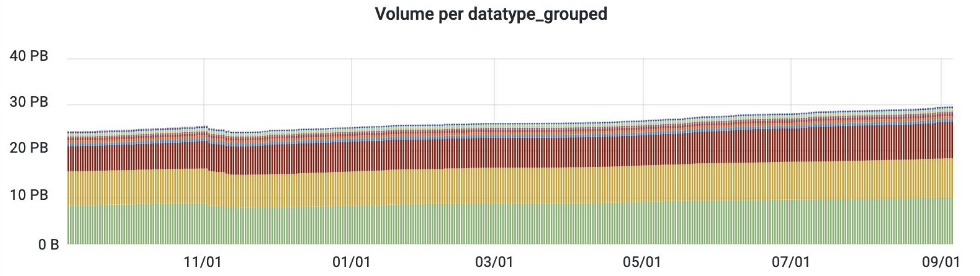

	min	max	avg	current
AOD	0 B	10.0 PB	8.87 PB	10.0 PB
HITS	0 B	8.44 PB	7.67 PB	8.44 PB
RAW	0 B	7.86 PB	6.56 PB	7.86 PB

28.13 PB occupato (31.77 PB disponibile)

Attività ATLAS ai T2 : CPU (settembre 2022-settembre 2023)

Referaggio 2023

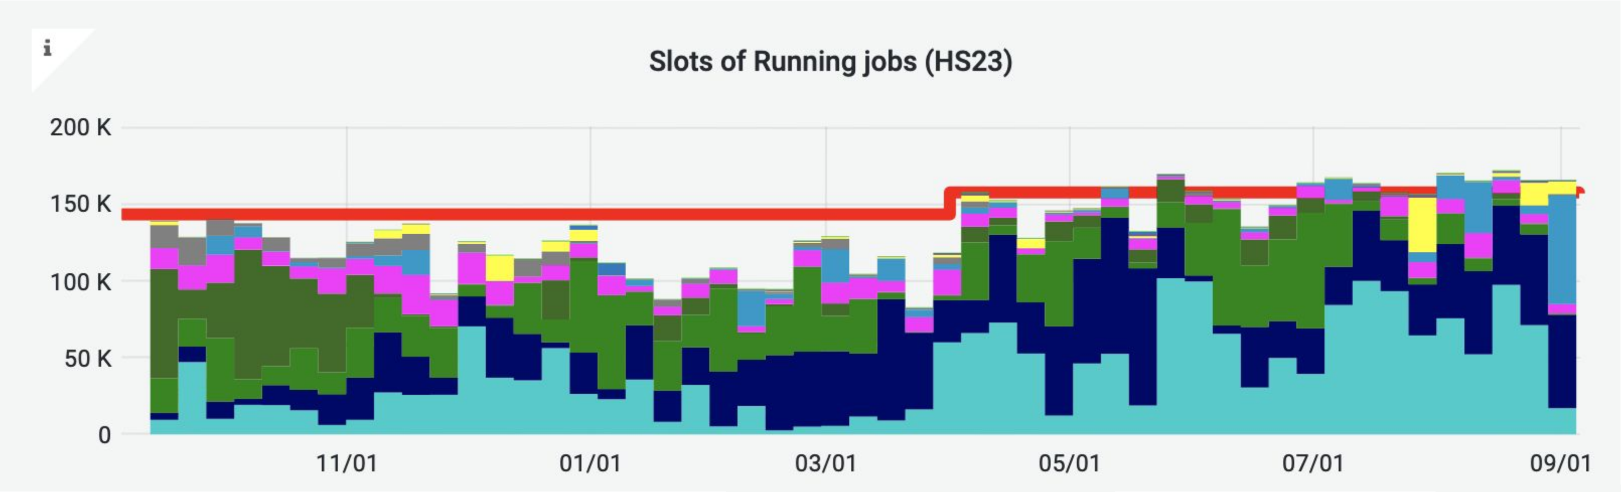

- Tipico share T2 :
  - MC (generation, simulation e reconstruction)
  - produzione di dati secondari (group production) e analisi utenti

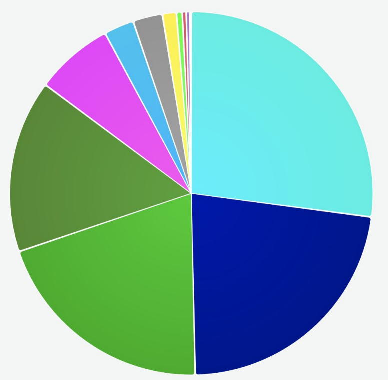

	Value	Percent
MC Simulation Full	1.06 Tri	27%
MC Reconstruction	883 Bil	23%
Group Production	787 Bil	20%
MC Event Generation	600 Bil	15%
User Analysis	270 Bil	7%
MC Simulation Fast	102 Bil	3%
Group Analysis	100 Bil	3%
Data Processing	45.6 Bil	1%

Attività ATLAS ai T2 : CPU

- In generale buon comportamento dei T2 italiani per l'ultimo anno di attività
  - reliability/availability ~ 95 %
  - Efficienza ~ 95%
  - Valori in linea con i migliori siti ATLAS

- Leggero under-pledge (~12%) sulla CPU fornita (average pledge 149 kHS23, computing power fornita 132 kHS23)
  - problemi di condizionamento in alcuni siti

- Tutti i siti sono ora up&running :
  - Ottima performance del sito di Milano che sta lavorando su tutti i workflows (ammodernamento apparati di rete)

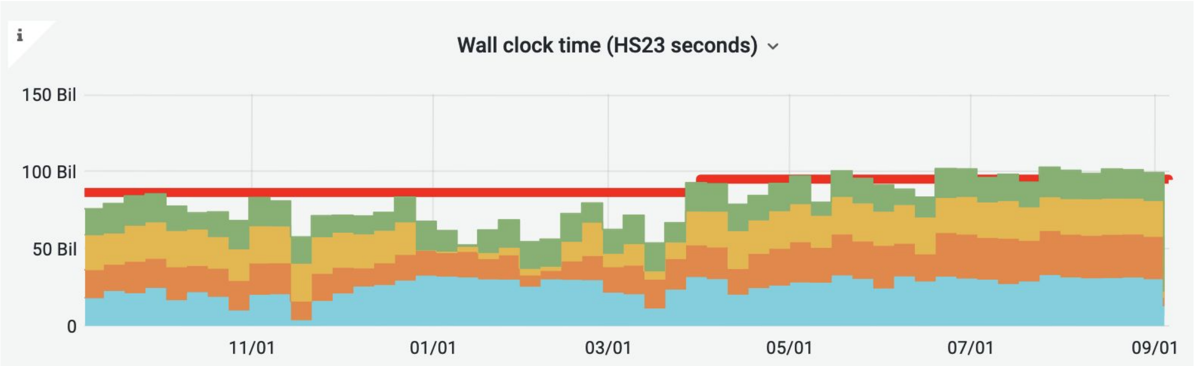

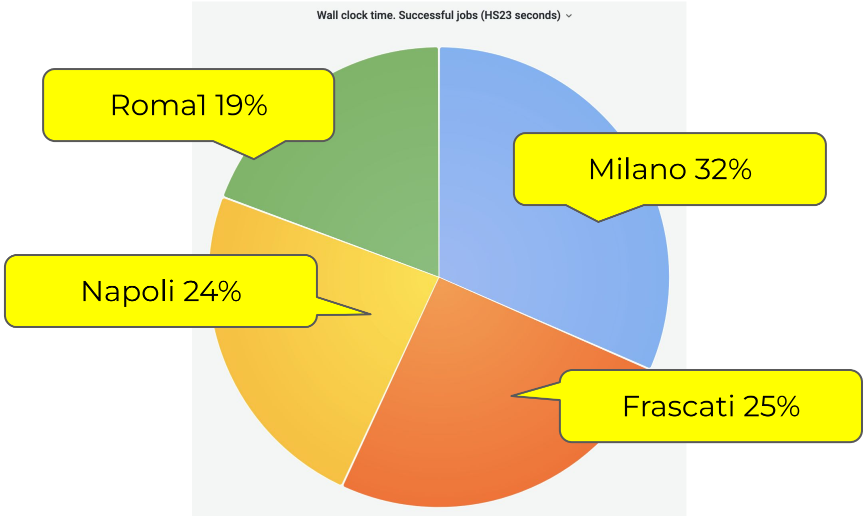

CPU performance summary (all sites)

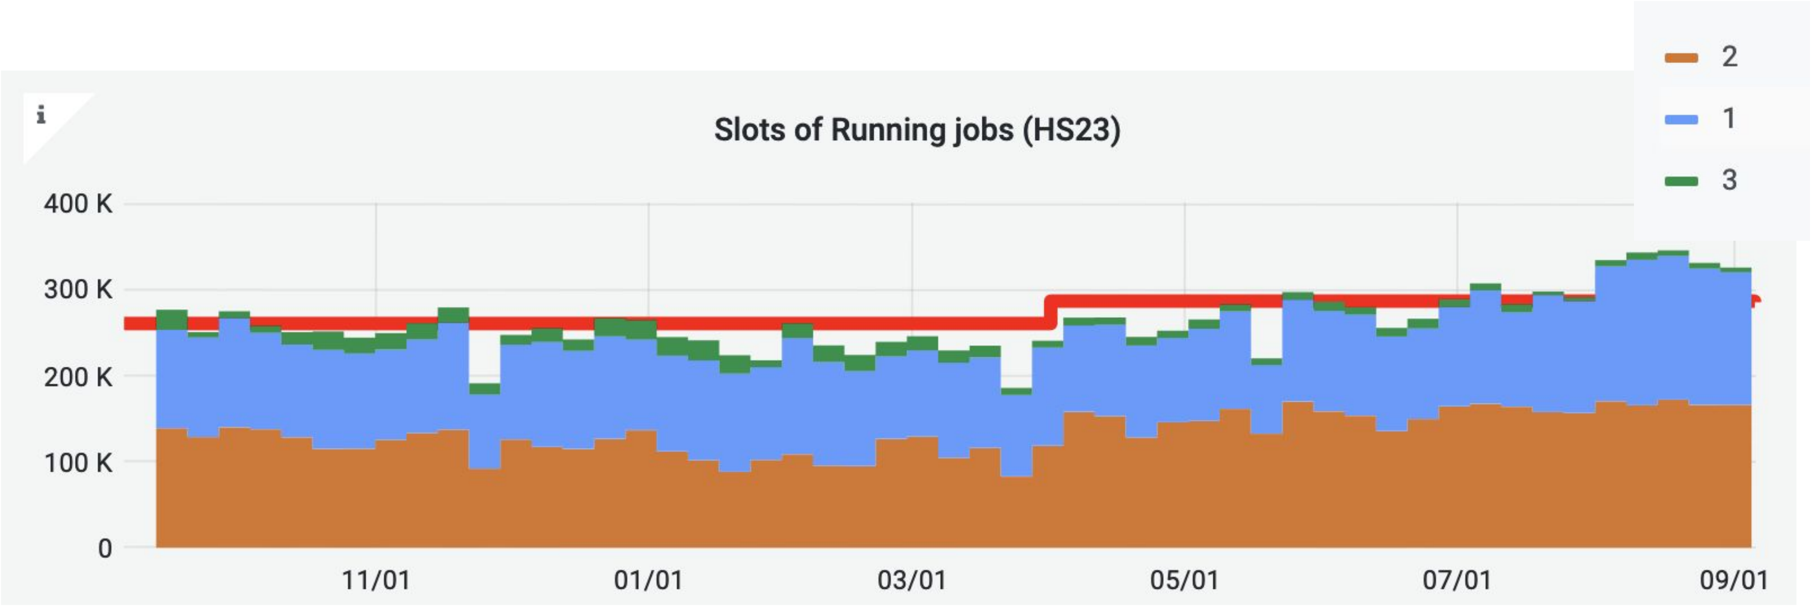

Globalmente leggera sofferenza dei siti italiani

- 271 kHS23 pledge, 260 kHS23 delivered: T2 e T1, contributo dei T3 (12 kHS23, ~ 5%) leggermente sotto la norma.
- Overall efficiency ~ 95% (in linea con l'efficienza della media dei siti ATLAS)

Situazione e disco e TAPE

	Disco occupato (PB) (dati ottobre 2023)	Disco allocato (PB) (dati ottobre 2023)	Teorici
Frascati	1.76	2.02	3.15
Milano	1.67	1.98	2.3
Napoli	4.19	4.58	4.15
Roma	1.54	1.75	2.45
tot.	9.16	10.33	12.05

NON validati

- Pledge 2022 : 9.94 PB
- Pledge 2023 : 11.76 PB

validati

	Disco occupato (PB) (dati ottobre 2023)	Disco allocato (PB) (dati ottobre 2023)	Tape (PB) (dati ottobre 2023)
CNAF	10.36	10.96	31.77

- Pledge 2022 (disco): 10.44 PB
- Pledge 2023 (disco): 12.24 PB
- Pledge 2022 (tape): 24.48 PB
- Pledge 2023 (tape): 31.7 PB

Contributo italiano (settembre 2022 - settembre 2023)

Contributo italiano alle attività di computing di ATLAS :

- ❏ CPU : overall ~ 6% wallclock time totale in HS23-sec (T1+T2+T3, solo risorse grid, T0 escluso)
  - ❏ leggermente sotto il contributo nominale per via del consistente overpledge di alcune cloud + contributo delle risorse HPC
- ❏ Disco occupato (datadisk, calibdisk, groupdisk e localgroupdisk) : overall ~6% (escludendo T0)
  - ❏ leggermente sotto share alcune risorse ancora da mettere il linea
- ❏ Tape (T1 only) : ~9%
  - ❏ perfettamente in linea con lo share (9%)

https://agenda.infn.it/event/34683/contributions/197356/attachments/105504/148335/20230523-Workshop-Loano.pdf

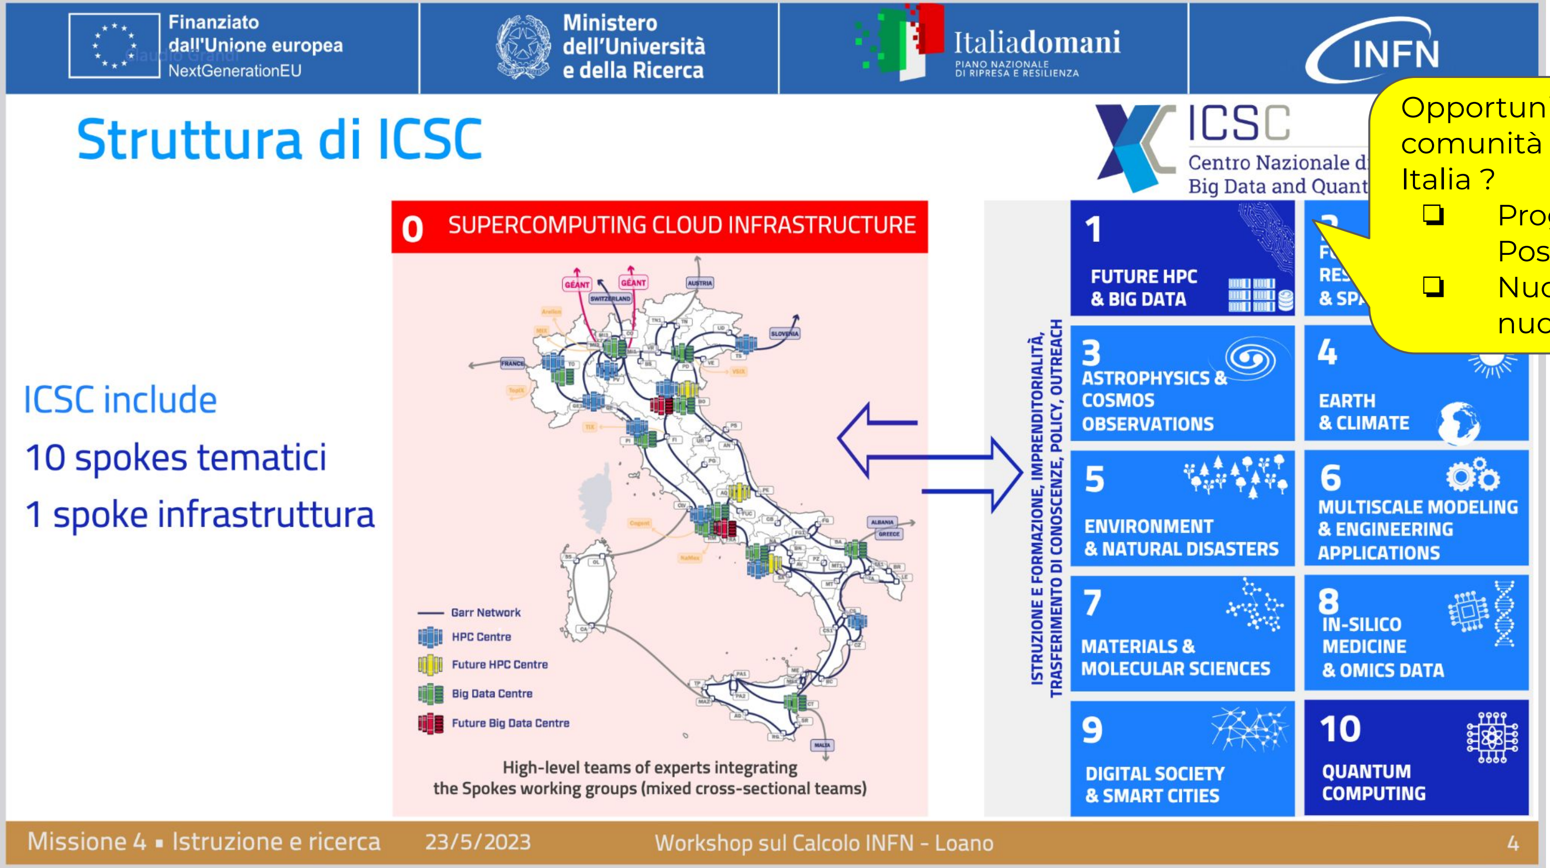

Opportunita' per la comunità di analisi di ATLAS Italia ?

- Progetto INCANT, Poster C. Marcon
- Nuovo personale, nuovi progetti di R&D

Richieste 2024 da RRB

CERN-RRB-2023-012

ATLAS		2022 C-RSG recomm.	2022 Pledged	2022 Used	2023 C-RSG recomm.	2023 Pledged	2024 Request	2024 req. /2023 C-RSG	2024 C-RSG recomm.
CPU	Tier-0	740	544	512	740	740	936	126%	936
	Tier-1	1300	1349	1545	1430	1520	1516	106%	1516
	Tier-2	1588	1616	3151	1747	1841	1852	106%	1852
	HLT	n/a	n/a	311	n/a	n/a	n/a	n/a	n/a
	Total	**3628**	**3509**	**5519**	**3917**	**4101**	**4304**	**110%**	**4304**
	Others			2161					
Disk	Tier-0	32.0	32.0	33.0	40.0	40.0	49.0	123%	49.0
	Tier-1	116.0	130.0	139.0	136.0	150.5	163.0	120%	163.0
	Tier-2	142.0	142.0	141.0	168.0	160.0	200.0	119%	200.0
	Total	**290.0**	**304.0**	**313.0**	**344.0**	**350.5**	**412.0**	**120%**	**412.0**
Tape	Tier-0	120.0	120.0	102.0	174.0	174.0	207.0	119%	207.0
	Tier-1	272.0	280.0	281.0	353.0	360.3	452.0	128%	452.0
	Total	**392.0**	**400.0**	**383.0**	**527.0**	**534.3**	**659.0**	**125%**	**659.0**

Richieste T1 2024

Referaggio 2023

ATLAS C-RSC recommended 2024

	ATLAS 2024 pledge	Italy share 2024	ATLASIT 2024 pledge	Richieste 2024	Richieste 2024 (k€)	ATLAS 2023 pledge	Italy share 2023	ATLASIT 2023 pledge
CPU T1 (HS06)	1516000	0,09	136440	7740	77400	1430000	0,09	128700
Disco T1 (TB)	163000	0,09	14670	2430	291600	136000	0,09	12240
Tape T1 (TB)	452000	0,09	40680	8910	89100	353000	0,09	31770
Totale (k€)					458100			

- Richieste calcolate rispetto al pledge del T1 (non rispetto alle risorse effettive attualmente disponibili)
- Mantenuto lo stesso share per ATLAS-IT (9%)
- Obsolescenza non inclusa
- Fattori di conversione : 120 €/TBn (disco), 10 €/HS23 (cpu), 10 €/TBn (tape)

Richieste T2 2024 complessive

Referaggio 2023

ATLAS C-RSG recommended 2024

	ATLAS 2024 pledge	Italy share 2024	ATLASIT 2024 pledge	ATLAS 2023 pledge	Italy share 2023	ATLASIT 2023 pledge
CPU T2 (HS06)	1852000	0,09	166680	1747000	0,09	157230
Disco T2 (TB)	200000	0,07	14000	168000	0,07	11760

- ❏ Fattori di conversione : 120 €/TBn (disco), 10 €/HS23 (cpu), 10 €/TBn (tape)
- ❏ Mantenuto lo stesso share ATLAS-IT (9% CPU e 7% disco)

Richieste dettagliate T2 : CPU

Referaggio 2023

CPU (HS06)	Frascati	Milano	Napoli	Roma1	Totale	Pledged
Totale CPU nei siti per il 2023 (HS06)	38312	39708	41886	40028	159934	157230
Obsolescenza 2024 (HS06)	10004 (0/10004)	5878 (5878/0)	0	0	15882 (5878/10004)	
Delta 2024 (HS06)	2500	2500	2500	2500	10000	
Totale richieste 2024 (HS06)	12504	8378	2500	2500	25882	
Totale CPU nei siti per il 2024 (HS06)	40812	42208	44386	42528	169934	166680

Previsione spesa (€)	125040	83780	25000	25000	258820

Richieste dettagliate T2 : disco

Referaggio 2023

DISCO (TB)	Frascati	Milano	Napoli	Roma1	Totale	Pledged
Totale disco nei siti per il 2023 (TB)	3148	2309	4147	2444	12048	11760
Obsolescenza 2024 (TB)	0	789	0	864	1653	
Delta 2024 (TB)	500	500	500	500	2000	
Totale richiesta 2024 (TB)	500	1289	500	1364	3653	
Totale disco nei siti per il 2024 (TB)	3648	2809	4647	2944	14048	14000

Previsione spesa (€)	60000	154680	60000	163680	438360

Hot topics : ATLAS Computing requirements for HL-LHC

ATLAS HL-LHC Computing Conceptual Design Report : projections of ATLAS computing requirements for Run3 and HL-LHC to fully exploit the machine physics potential is quite scaring !

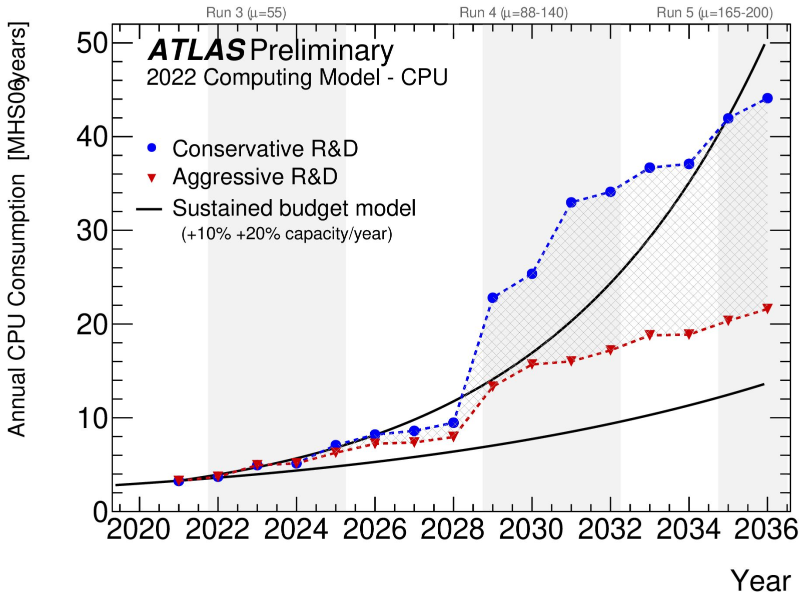

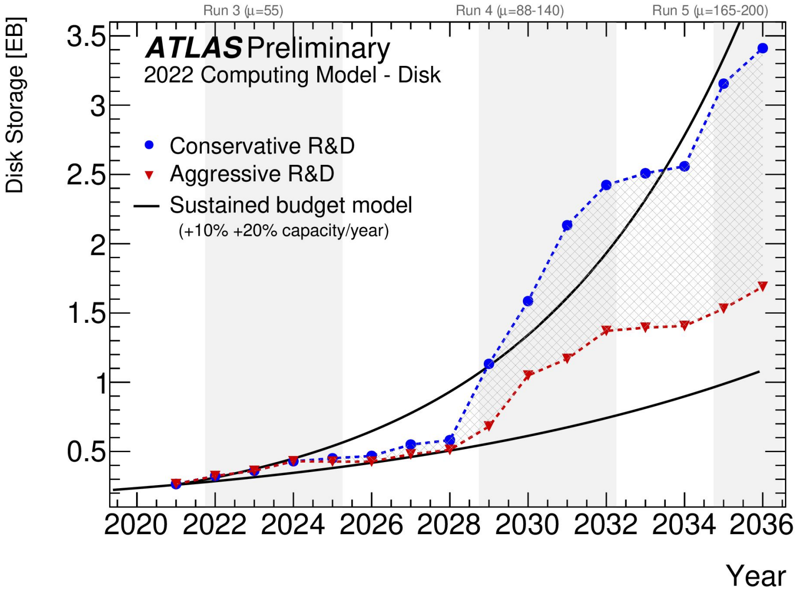

Discussion started on possible strategies to meet the demanding requirements of HL-LHC computing

- optimisation (both speed and flexibility) of the experiment (e.g. reconstruction, simulation) and non-experiment (e.g. generation) software
- optimisation of the available hardware infrastructure usage

Sustainability in HEP computing

https://agenda.infn.it/event/34738/contributions/191871/attachments/106878/150865/SustainComputingLohwasser_v2.pdf

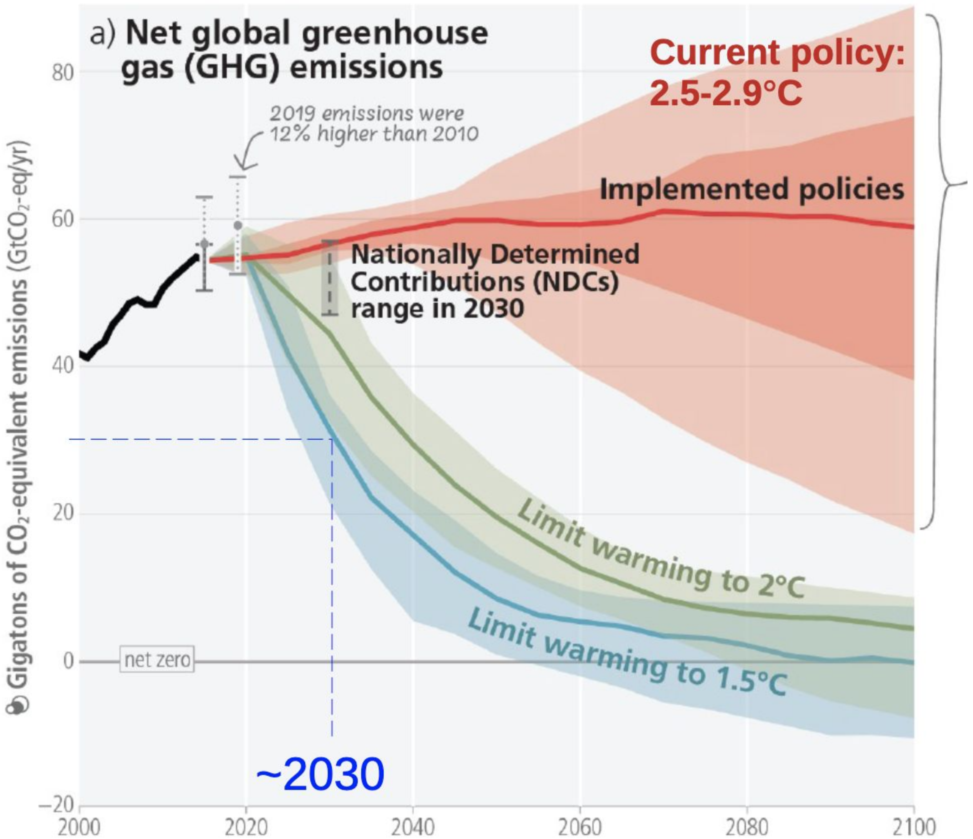

- ❏ Reduction to zero emissions around 2100 : 50% of the reduction should be achieved by ~2030 → in 7 years
- ❏ WLCG requirements : from ~ 100% increase in energy needs (pessimistic scenario) to 10% (optimistic scenario)

How to reconcile these ?

Sustainability in HEP computing

https://agenda.infn.it/event/34738/contributions/191871/attachments/106878/150865/SustainComputingLohwasser_v2.pdf

https://indico.cern.ch/event/1211083/contributions/5225666/attachments/2580240/4450193/2023-CERN-ATLAS-Geneva-final.pdf

- Expand CO_2-free energies (factor 12)
  - Renewable power for computing: processors and cooling;
  - Consider district heating and site selection;
  - Job scheduling according to energy availability; ...
- Increase energy efficiency (factor 2)
  - Optimised processors (clocks, GPUs),
  - architecture, cooling system,
  - software, quantum computing?, ...
- Save energy (factor 2)
  - Prioritise research questions
  - Optimise debugging, statistics and precision;
  - Modular and reusable software;
  - Modular and repairable hardware, reduce purchases;

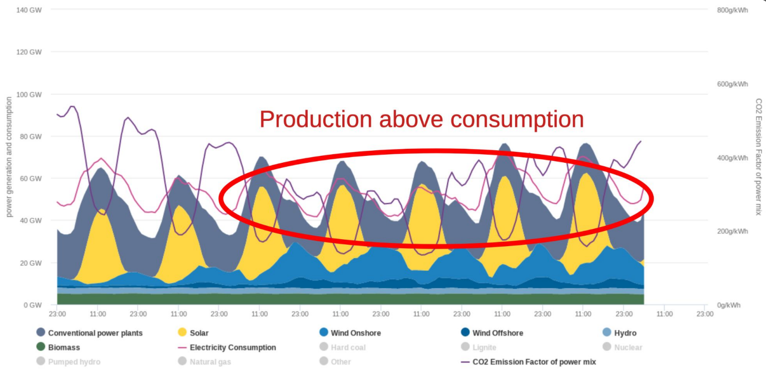

Rather than having the grid adapt to the data center, we need to have the data center adapt to the grid .. The data center must “dance with the grid.”

Reminder: <u>Paris agreement is in principle legally binding</u>
- pressure on us / our savings might need to be increased
- gives us negotiating power if we have a clear plan and strategy with demonstrable impacts and realistically achievable objectives in line with 1.5°C

Sustainability in HEP computing

https://agenda.infn.it/event/34738/contributions/191871/attachments/106878/150865/SustainComputingLohwasser_v2.pdf

- ❏ Expand CO2-free energies (factor 12)
  - ❏ Renewable power for computing: processors and cooling;
  - ❏ Consider district heating and site selection;
  - ❏ Job scheduling according to energy availability; ...
- ❏ Increase energy efficiency (factor 2)

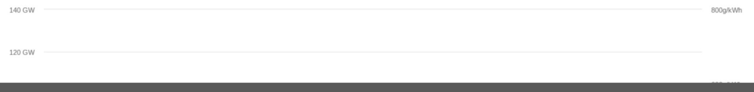

Estimated Carbon Footprint for the Task

Category	gCO2
Succeeded	56.09 gCO2
Failed	0 gCO2
Cancelled	0 gCO2
Total	**56.09 gCO2**

  - ❏ Modular and reusable software;
  - ❏ Modular and repairable hardware, reduce purchases;

Rather than having the grid adapt to the data center, we need to have the data center adapt to the grid .. The data center must “dance with the grid.”

Reminder: Paris agreement is in principle legally binding
- ❏ pressure on us / our savings might need to be increased
- ❏ gives us negotiating power if we have a clear plan and strategy with demonstrable impacts and realistically achievable objectives in line with 1.5°C

More efficient full simulation

https://indico.cern.ch/event/1180455/contributions/5318122/attachments/2672908/4633983/AtlasWeek_june23.pdf

- Several improvements in the CPU performances of full simulation (with no loss of accuracy) have been implemented in the last two years : full simulation is about 2 times faster than it was for run 2.
- Work to further speed up full simulation for future MC campaigns and HL-LHC continue
- A new EMEC geometry (using G4 supported shapes) makes the simulation of EMEC 2.5 times faster and GPU compatible
- In parallel, prototypes of GPU simulation are being developed in collaboration with Adept and Celeritas (two groups developing GPU simulation across experiments).
- Expect to be able to test a GPU workflow for the HL LHC Computing TDR next year

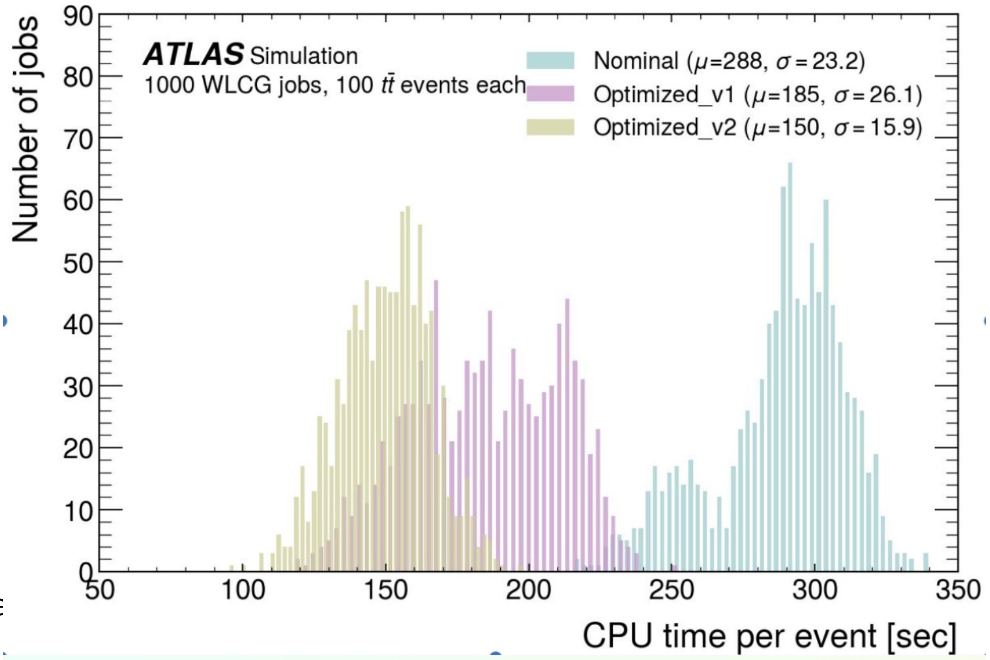

Increase the usage of Fast Simulation

- ❏ Use parameterizations of calorimeter response (FastCaloSim V2) and a generative NN approach (FastCaloGAN)
- ❏ Fast CaloSim parameterization and FastCaloGAN trainings have been done to match the G4 version (10.6) used for run 3
- ❏ FastCaloGAN developments results in better agreement with Geant4 than the run 2 version, also used for more particles and energies
- ❏ Fatras : Fatras replaces Full ID simulation (dominant CPU consumer in AF3) with a simple tracking geometry and parameterizations of interactions

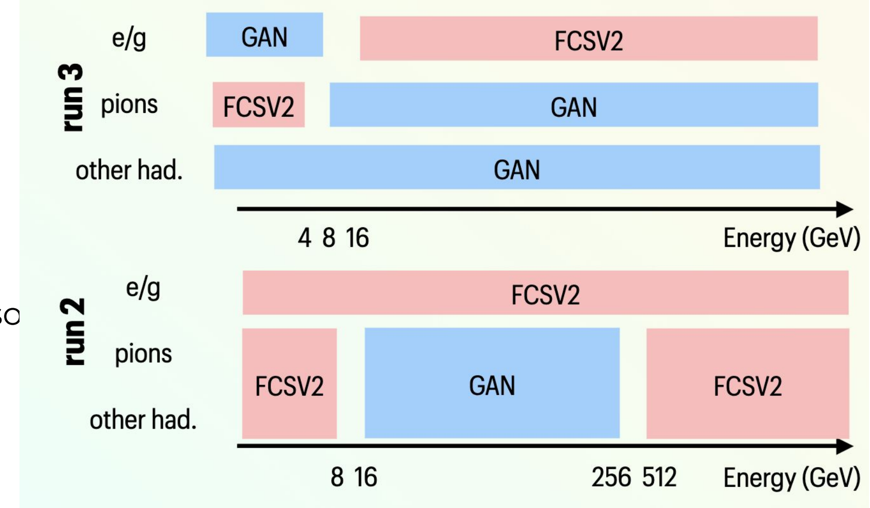

- ❏ Track overlay : MC overlay adds pre-digitized pileup at raw data level, recycling the same pileup events for different hard scatter events
  - ❏ Track overlay also uses pre-reconstructed pileup tracks, saving the CPU for pileup track reconstruction

Heterogeneous computing

https://indico.cern.ch/event/1180455/contributions/4958179/attachments/2675554/4639712/Heterogeneous%20Computing%20in%20ATLAS.pdf

Why Use Accelerator Resources?

BERKELEY LAB

- Because they're here / there
  - Large fraction of computing power comes from GPUs in most HPCs and increasing number of HTCs
  - lxplus-gpu, grid nodes, swan instances, cloud, under your desk and more
- Because they're energy efficient
  - Computing sites are increasingly limited by the amount of power they consume
  - GPUs are more power efficient than CPUs per FLOP
- Because we always need more computing power
  - But is it cheaper or more efficient to get there with GPUs or FPGAs instead of traditional CPUs?

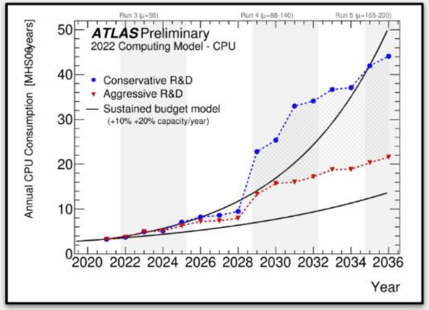

- Because everyone else is doing it
  - ALICE, LHCb, CMS, ChatGPT, cryptobros ....

Heterogeneous computing

https://indico.cern.ch/event/1180455/contributions/4958179/attachments/2675554/4639712/Heterogeneous%20Computing%20in%20ATLAS.pdf

BERKELEY LAB

Summary

ATLAS has many different avenues exploring the use of accelerators for both the online EF Tracking for offline reconstruction and analysis.

We have not yet decided whether we **NEED** GPUs/FPGAs for TDAQ / EF.

- technology decision will be in 2025
- many results of this investigation will be leveraged by offline software to take advantage of existing GPU resources (grid, cloud, HPC, laptop, etc)
  - we **WILL** use GPUs in offline reco and analysis

Calorimeter simulation, EMCalo clustering, analysis optimisations

Usage of Machine Learning techniques in both online and offline continues to increase.

- can experience tremendous benefits from accelerators, especially in training of models
- working on finding efficient ways of integrating ML as a service into our workflows

Heterogeneous computing

- ❑ ARM CPUs have low cost and minimal power consumption.
- ❑ Some future HPCs plan to have ARM CPU partitions
- ❑ Athena, AthSimulation, AnalysisBase ported to the aarch64-centos7-gcc11-opt platform last year
- ❑ Technical physics validation successfully passed using Amazon
- ❑ Cloud Gravition2 ARM CPU fully integrated into PanDA/Rucio:
  - ❑ Full simulation Athena,23.0.3 (ATLPHYSVAL-872)
  - ❑ Reconstruction Athena,23.0.14 (ATLPHYSVAL-919)
- ❑ 3 ATLAS workflows (Sherpa, Geant4, Reco) integrated into the new WLCG HEPscore benchmark available for x86_64 and ARM
- ❑ WLCG sites start to get interested in ARM
  - ❑ Glasgow team shows competitive result for events/kWh and throughput

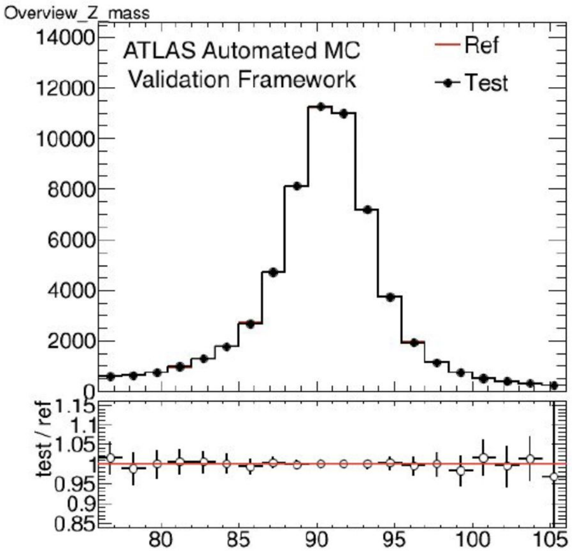

Reconstructed MC $Z\rightarrow\mu\mu$ invariant mass

Analysis model(s)

https://indico.cern.ch/event/1180455/contributions/5359601/attachments/2673977/4637700/AMG-2023-06-27.pdf

- The output of the data and MC reconstruction is stored in Analysis Object Data (AOD) files
- In run 2 these datasets were processed in the derivation framework which produces about 80 different derived AOD (DAOD) formats .
- In run 3 DAOD_PHYS now the default format for analysis
- One more DAOD format (especially in view of run4/5) : PHYSLITE
  - object CP corrections already applied
  - extra good object thinning
  - not all content from PHYS
- future: use PHYSLITE like an n-tuple push towards columnar approach + CP tools, closer to python ecosystem (lot of work ongoing on this subject)
- future:
  - event augmentation : save extra data for some events, further reduce need for custom DAODs
  - lossy compression

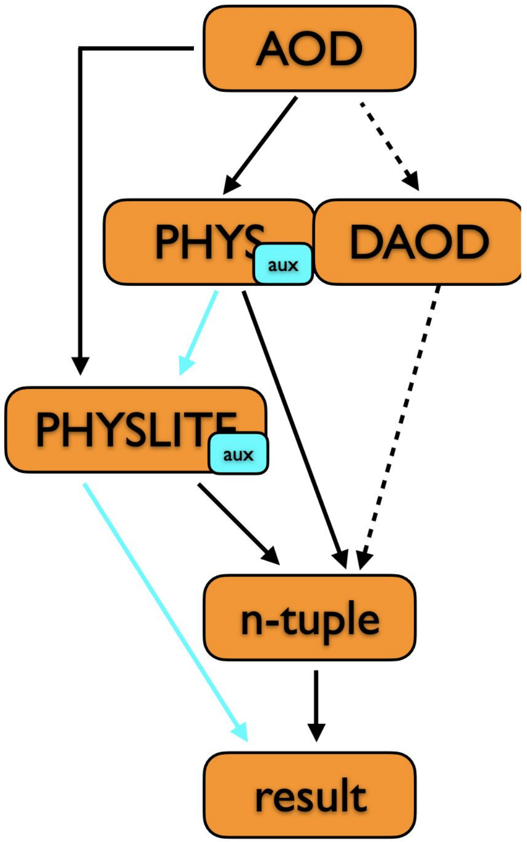

Lossy compression

https://indico.cern.ch/event/1295808/contributions/5446232/attachments/2665301/4618464/DAODPHYS_LossyCompression.pdf

- In order to strengthen the specific characteristics of the PHYSLITE format, studies are ongong to :
  - study the minimum necessary variable precision to reduce the size after compression;
  - review and remove unnecessary content.
- Lossy float compression on DAOD_PHYSLITE can result in a significant reduction of the storage footprint by removing unnecessary bytes;

Original file [M]	48	-
High compressed file [M] *	34	- ~30%
Low compressed file [M] *	41	- 15%

- All these steps must be implemented very carefully to prevent loss of data information -> they must not have an impact on the final physics results and on the uncertainties;
- This is beneficial as long as the additional source of error introduced by the truncation/compression is not dominant over the observables' instrumental precision;
- Interactions with performance/analysis groups ongoing in order to identify this optimal limit.
- (more info here https://cds.cern.ch/record/2865166?ln=it for non lossy compression studies)

https://cds.cern.ch/record/2870346/files/ATL-SOFT-PROC-2023-027.pdf

- The ATLAS experiment collects several billion of collisions from the LHC every year. In addition, simulated events are generated to compare the results of the analysis of real data with different physics models.
- All these data need to be catalogued in a large, high performance and high reliability system that can provide ATLAS members the possibility to search for and retrieve one or more individual events from the tens of millions of data files, in order to perform detailed checks, more refined analyses or produce event displays – the so-called “event picking” operations
- The EventIndex design started in 2013 and the first implementation was fully functional for the start of LHC Run 2 in 2015 : necessary to upgrade it in order to stand the expected higher data rates for Run 3 (2022 onwards) and beyond.
- The new EventIndex for LHC Run 3 have demonstrated improved efficacy in handling large amounts of data and accessibility for users.
- In march 2023 the data store contained 540 billion event records belonging to 280’000 datasets, occupying 47 TB in HBase.

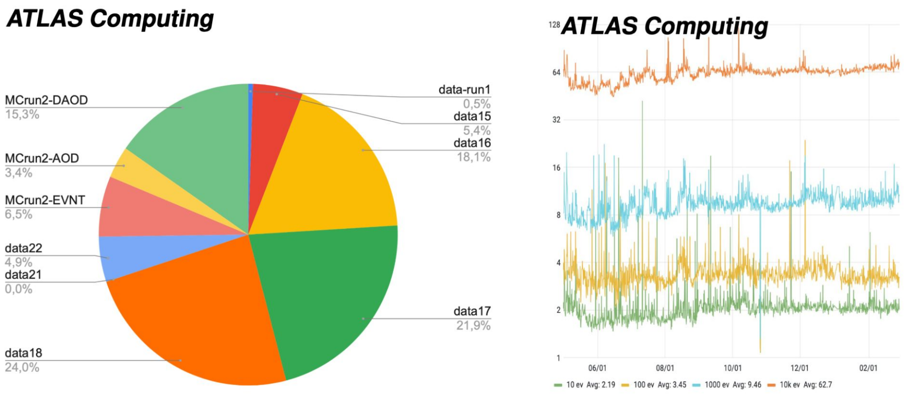

Figure 3. Left: fraction of event records stored by data type at the end of March 2023. Right: lookup times in seconds for queries for 10, 100, 1000 and 10’000 events from performance tests executed between May 2022 and February 2023.

Un paio di punti dall'ultimo ICB

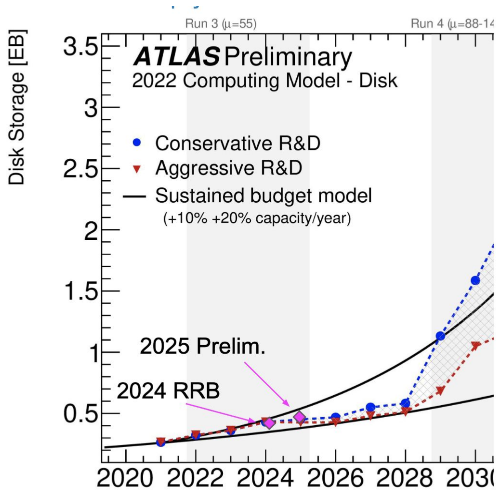

2025 Prelim.: **469 PB**
Tier0 contrib.: 56 PB
WLCG contrib.: 410 PB

2028: 550–610 PB
2029: 680–1220 PB
2030: 1050–1630 PB

23–100% increase in 2029
55–34% increase in 2030

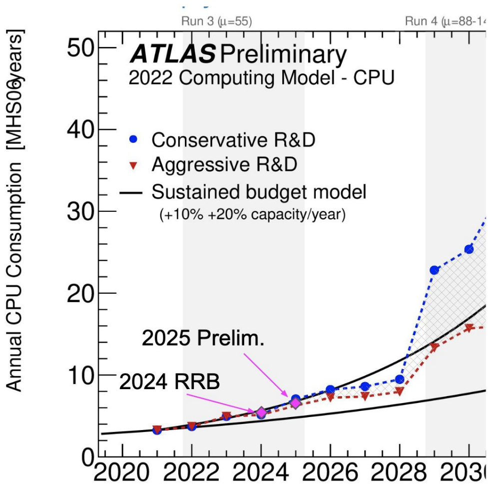

Current biggest change is from 2028 to 2029

2028: 7.8 – 8.4 MHS23
2029: 11.8 – 21.2 MHS23

Increase is 50–150%; roughly the same as the integrated change from 2029 to 2036

With this large single-year increase coming, **how can we formulate our requests** to best help you ensure that you are able to deliver resources in the HL-LHC period?

- The "big jump" nel 2028 : sono necessarie azioni di 'preavviso' verso le funding agencies ? Come procedere per mitigare l'impatto ?
- Grande interesse verso Leonardo : possibile uso opportunistico delle risorse a la VEGA ?
- Molto interesse verso i nuovi siti ARM (tests anche al CNAF)
- ADC : il personale esperto diminuisce, più' persone che lasciano rispetto alle nuove cha arrivano

Conclusioni

- ATLAS ha fatto girare in modo smooth 650k jobs slot in parallelo
- I siti italiani hanno continuato a lavorare in modo efficiente per tutto il periodo settembre 2022-settembre 2023. Siti affidabili (elevata reliability e availability) e con efficienza paragonabile ai migliori centri ATLAS
- Presentata la situazione delle risorse disponibili (gare) e le richieste per il 2024
  - Mantenuti gli stessi share delle scorse richieste
- Enorme sforzo per preparare la roadmap per soddisfare gli esigenti computing requirements per HL-LHC e sostenibilità: affrontare il problema da diverse prospettive, molte attività di R&D in corso
- Carenza di manpower segnalata molte volte : disponibilità ad incrementare il nostro coinvolgimento ? Anche per migliorare l'in-kind contribution !
- Menu di oggi :
  - Stato installazione nuovo hardware
  - Stato accounting, confronto tra i siti
  - Stato e programma upgrades infrastrutturali
  - Migrazione storage : piani di Roma1 e Frascati. Nuovo storage dCache anche a Milano ?
  - Readiness for DC24
  - Piani upgrade condor e linux
  - Interazione con i WP di ICSC

Backup

Richieste finanziarie T2

	Frascati	Milano	Napoli	Roma1	Totale
Spesa CPU (€)	125040	83780	25000	25000	258820
Spesa DISCO (€)	60000	154680	60000	163680	438360
Totale (€)	185040	238460	85000	188680	697180
Overhead server (CSN1)	12952	16692	5950	13207	48802
Overhead rete (DATACLOUD)	10502	12760	4500	9684	37447
Totale (overheads inclusi) (€)	208494	267912	95450	211571	783429

- Overheads :
  - quota server (su CSN1): 7% del totale CPU + disco
  - quota rete (su Datacloud): 6% CPU + 5% disco
- Fattori di conversione : 120 €/TBn (disco), 10 €/HS06 (cpu), 10 €/TBn (tape)

Richieste dettagliate T2 : CPU

CPU (HS06)	Frascati	Milano	Napoli	Roma1	Totale	Pledged
Totale CPU nei siti per il 2023 (HS06)	38312	39708	41886	40028	159934	157230
Obsolescenza 2024 (HS06)	10004 (0/10004)	5878 (5878/0)	0	0	15882 (5878/10004)	
Delta 2024 (HS06)	2500	2500	2500	2500	10000	
Totale richieste 2024 (HS06)	12504	8378	2500	2500	25882	
Totale CPU nei siti per il 2024 (HS06)	40812	42208	44386	42528	169934	166680

Previsione spesa (€)	125040	83780	25000	25000	258820

Totale risorse PNRR da acquisire nel 2023 (HS06)	51194	51194	121624	25605	249617	
Pledge 2023 promesso ad ATLAS	19866	11426	3000	17521	51813	
Risorse PNRR 2023 non ATLAS (HS06)	31328	39768	118624	8084		
Avanzo dopo copertura pledge 2024 ATLAS	18824	31390	116124	5584		

Richieste dettagliate T2 : disco

DISCO (TB)	Frascati	Milano	Napoli	Roma1	Totale	Pledged
Totale disco nei siti per il 2023 (TB)	3148	2309	4147	2444	12048	11760
Obsolescenza 2024 (TB)	0	789	0	864	1653	
Delta 2024 (TB)	500	500	500	500	2000	
Totale richiesta 2024 (TB)	500	1289	500	1364	3653	
Totale disco nei siti per il 2024 (TB)	3648	2809	4647	2944	14048	14000

Previsione spesa (€)	60000	154680	60000	163680	438360

Lordi : -30% netti

Totale risorse da acquisire PNRR 2023 (TB)	2379	2230	5655	1200	11464	
Pledge 2023 promesso ad ATLAS	0	0	0	0	0	
Risorse PNRR 2023 non ATLAS (TB)	2379	2230	5655	1200	11464	
Avanzo dopo copertura pledge 2024 ATLAS	1879	941	5155	-164		

Utilizzo risorse ATLAS (settembre 2022 - settembre 2023) : CPU

- Circa 650k (average) computing slots ogni giorno per tutto il periodo in esame (pledged e unpledged, tutti i tiers insieme)

- Intensa attivita' di <u>produzione MC</u> (~70% includendo generazione, simulazione e ricostruzione) e <u>analisi dati</u> (~30% tra derivations e analisi vera e propria)

- Efficienza (wallclock consumption kHS23 seconds) ~ 92%

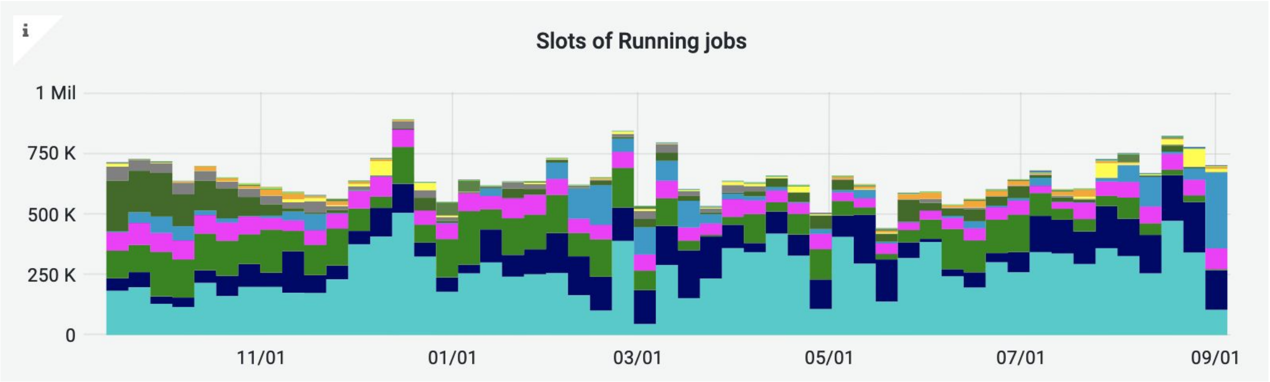

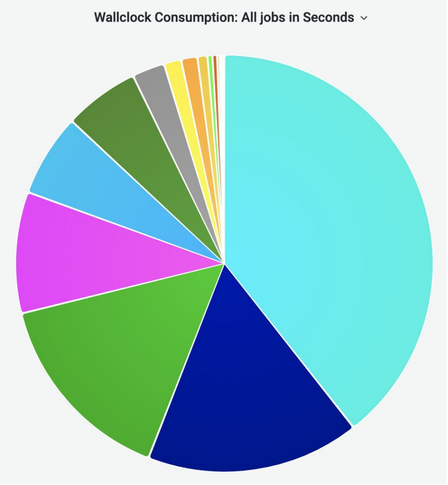

MC Simulation Full	7.97 Tri	40%
MC Reconstruction	3.34 Tri	17%
Group Production	3.07 Tri	15%
User Analysis	1.89 Tri	9%
MC Simulation Fast	1.28 Tri	6%
MC Event Generation	1.17 Tri	6%
Group Analysis	497 Bil	2%
Data Processing	256 Bil	1%
t0_processing	238 Bil	1%
MC Resimulation	142 Bil	1%

Utilizzo risorse ATLAS (settembre 2022 - settembre 2023) : CPU

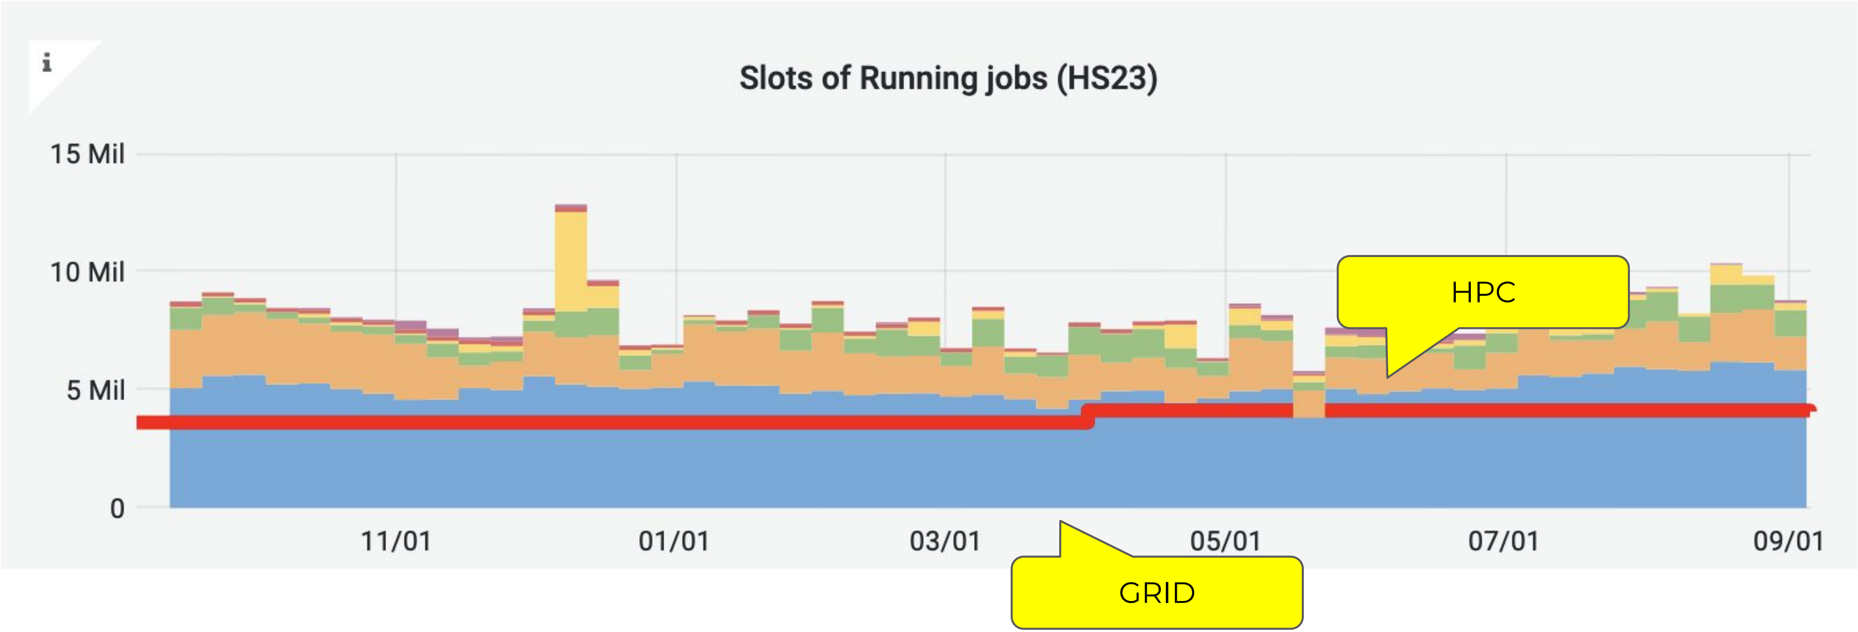

- Significativo over-pledge da siti grid (~ 30 % del pledge)
- Significativo contributo da risorse HPC opportunistiche (45 % del pledge, VEGA)

Utilizzo risorse ATLAS (settembre 2022 - settembre 2023) : disco

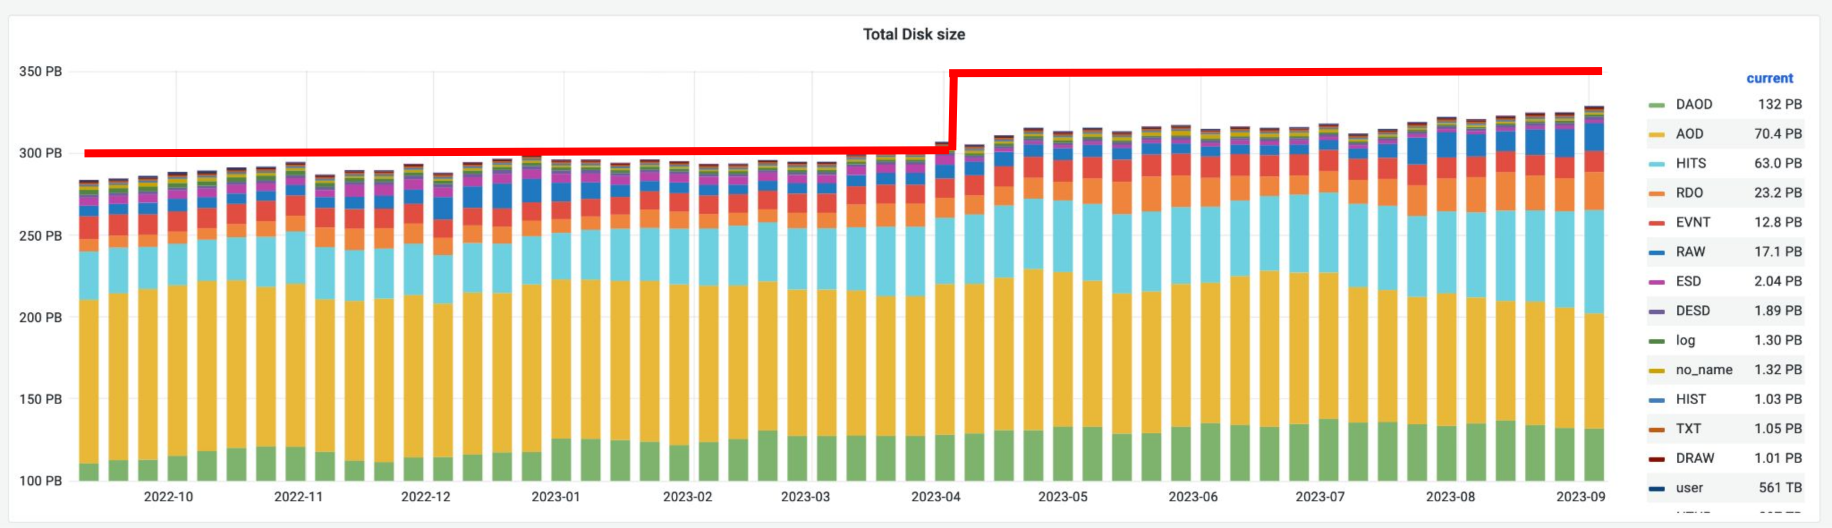

- ❑ Disco occupato ~ 329 PB (~700 M files) su 350 PB (pledge): principalmente dati derivati (DAOD)

Utilizzo risorse ATLAS (settembre 2022 - settembre 2023) : tape

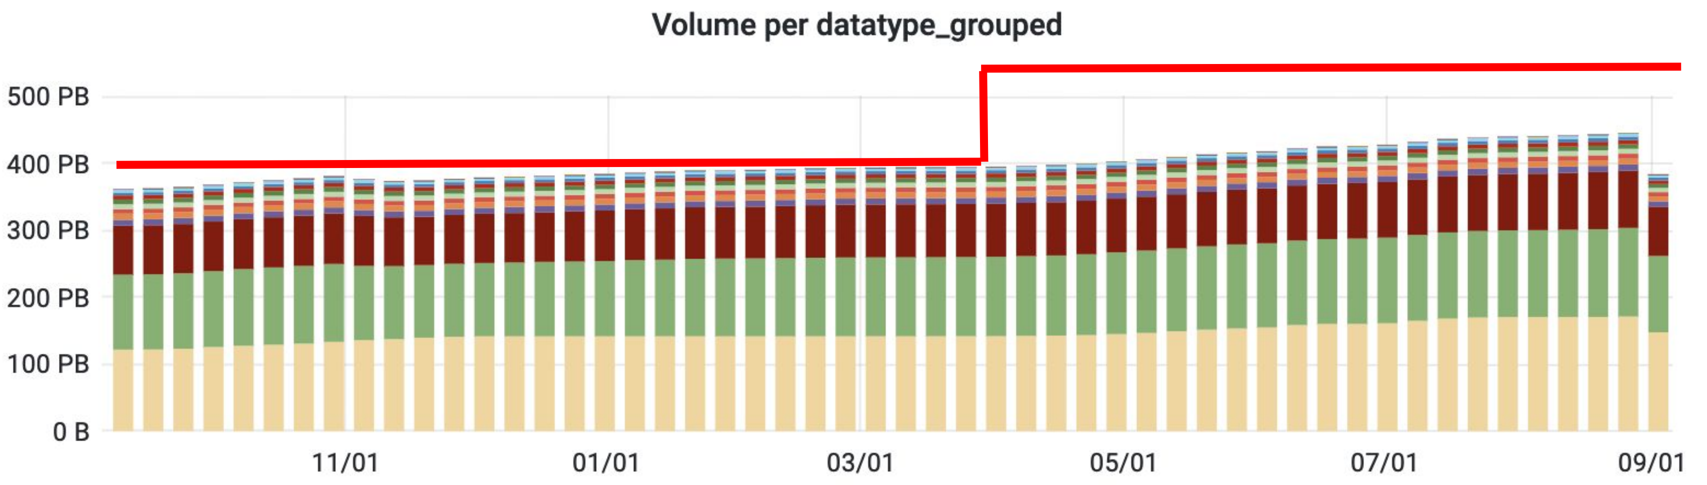

	min	max	avg	current
RAW	34.8 PB	171 PB	144 PB	148 PB
AOD	32.0 PB	132 PB	117 PB	114 PB
HITS	20.8 PB	85.6 PB	77.4 PB	73.8 PB

- Tape occupato ~ 448 PB (~300 M files) su 534 PB (pledge): mainly RAW data, AOD e HITS

Proiezione obsolescenze

Proiezione obsolescenze CPU (HS06)

	2023	2024	2025	2026	2027	2028
Frascati	16866	10004		8440		
Milano	8003	5877			14000	8400
Napoli	0		2667			36200
Roma	8003			14000		8400
totale	32872	15881	2667	22440	14000	53000

Proiezione obsolescenze disco (TBn)

	2023	2024	2025	2026	2027	2028
Frascati	0			561		1800
Milano	0	800		850		
Napoli	0			500		3500
Roma	0	800				
totale	0	1600	0	1911	0	5300

Multi-threaded reconstruction software (AthenaMT)

ATL-SOFT-PUB-2021-002

- The benchmark jobs use real data from run 357750 taken during 2018, with 250 events per worker process or thread.
- The data events have an average number of interactions per bunch crossing ⟨μ⟩ = 50, which is approximately that expected for the luminosity-leveling period during Run 3.
- Tests on an Intel®Xeon®CPU E5-2630 v3 at 2.40 GHz (16 cores no SMT) machine + 126 GB of memory.
- Multi-threaded reconstruction software allows a better exploitation of opportunistic resources (eg HPC)

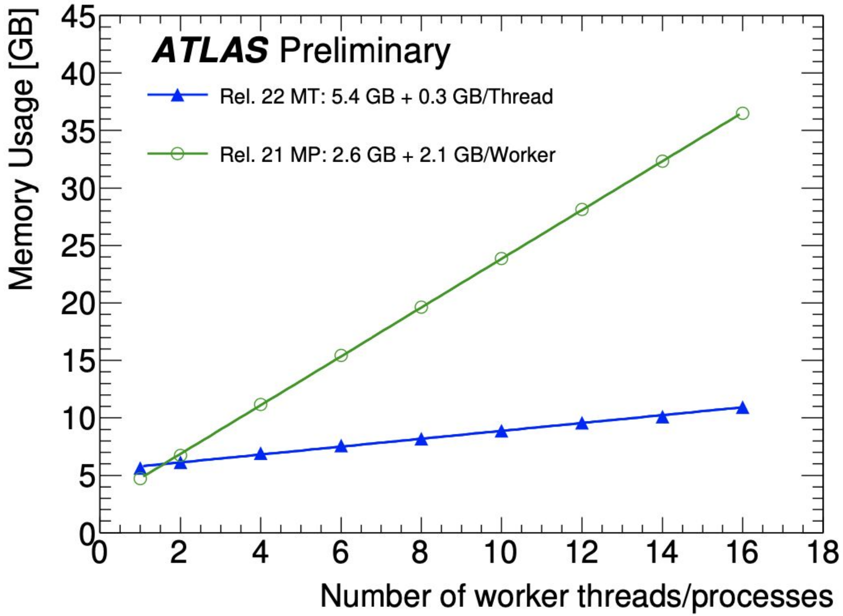

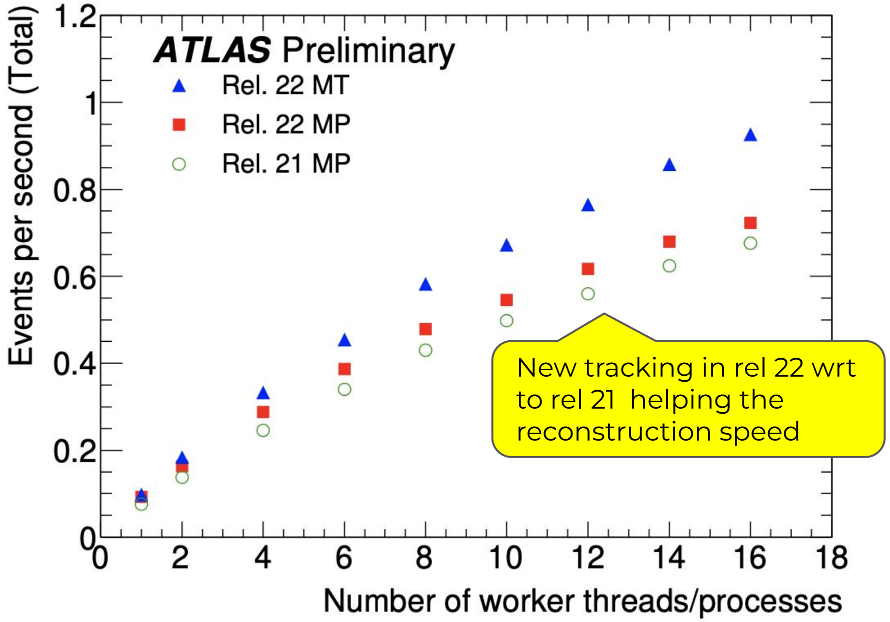

A new analysis model

The ATLAS Run-2 analysis model has been highly successful in the view of the productivity of ATLAS, but it has been expensive in terms of resource usage. The ATLAS Analysis Model Study Group for Run-3 (AMSG-R3) setup at the end of Run-2 was tasked to analyse the efficiency and suitability of the current model and to propose significant improvements.

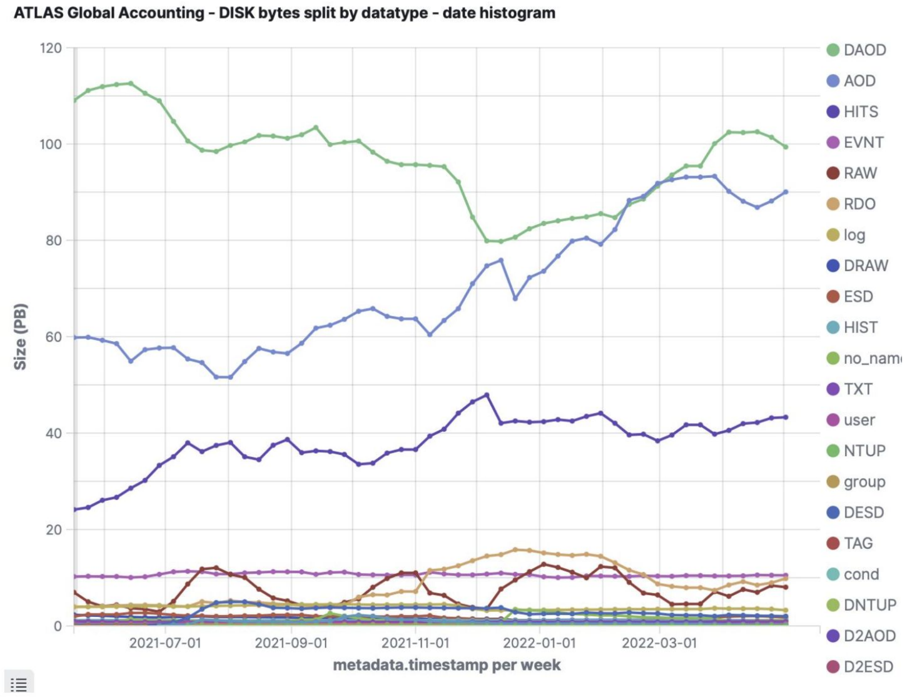

- The size/event for the AOD is about 600 kB
- Around 80 DAOD formats, size in the range of 40-450 kB depending on the type of the physics selection and the information retained.
- only 1-2 replicas of each dataset and campaign can be kept on disk.
- AODs and DAODs which are the two formats taking more than 70% of the disk space today
- As a rough Run-2 input parameter an initial sum of 132 PB of disk space used for AOD and DAOD format
- For the HL-LHC the projections of the ATLAS needs are significantly over the yearly flat budget increase. ATLAS is therefore investing significantly in methods to reduce the disk space needs in several areas